COLLEGES BRETAGNE LOIRE — SCIENCES POUR L'INGENIEUR ET LE NUMERIQUE

ENSTA — INSTITUT POLYTECHNIQUE DE PARIS

THÈSE DE DOCTORAT DE

L'ENSTA

ÉCOLE DOCTORALE N° 648
Sciences pour l'Ingénieur et le Numérique
Spécialité : *Génie industriel, productique, automatique et robotique ; Mécanique des solides, des matériaux, des structures et des surfaces*

Par
Mirado MORTEL

Natural Locomotion Families

Principle, Method, and Control

Soutenance prévue à Brest, le 25 novembre 2026 à 14 h 30
Unité de recherche : Lab-STICC – UMR CNRS 6285 – équipe ROBEX

Rapporteurs avant soutenance :

Paolo ROBUFFO GIORDANO	Directeur de recherche CNRS – HDR, CNRS / Rennes
Nicolas MANSARD	Directeur de recherche CNRS – HDR, CNRS / Toulouse

Composition du jury proposée :

Président	À préciser après la soutenance	
Rapporteur	Paolo ROBUFFO GIORDANO	Directeur de recherche CNRS – HDR, CNRS / Rennes
Rapporteur	Nicolas MANSARD	Directeur de recherche CNRS – HDR, CNRS / Toulouse
Examinatrice	Marie BABEL	Professeure des universités – HDR, INSA Rennes / Rennes
Examinateur	Frédéric BOYER	Professeur de l'Institut Mines-Télécom – HDR, IMT Atlantique / Nantes
Directeur de thèse	Luc JAULIN	Professeur des universités – HDR, ENSTA / Brest
Co-direction / co-encadrement	Lionel LAPIERRE	Professeur de l'ENSTA – HDR, ENSTA / Brest
Co-encadrement	Simon ROHOU	Maître de conférences, ENSTA / Brest

Version de lecture V7.1 – 24 septembre 2026. Jury proposé, sous réserve d'agrément.

Résumé

Pourquoi courir paraît-il plus naturel que marcher toujours plus vite ? Cette intuition invite à chercher, pour un robot, les démarches que sa mécanique favorise. Cette thèse cherche à les définir, les calculer et les entretenir. Elle considère d'abord un monde idéal sans dissipation ni actionnement, pour distinguer l'organisation du mouvement de l'énergie nécessaire à son entretien.

Les oscillations naturelles donnent un premier repère. Près d'un équilibre, un mode linéaire conserve sa fréquence et sa coordination lorsque l'énergie augmente. À amplitude finie, la fréquence du pendule simple varie déjà ; avec plusieurs articulations, la coordination peut changer aussi. Les modes non linéaires décrivent ces familles d'oscillations : le plan du mode linéaire devient une surface courbe dans l'espace d'état. Sur une surface modale régulière, l'intuition action–angle donne deux repères : l'action choisit le cycle, l'angle situe le mouvement sur ce cycle.

Osciller ne suffit pourtant pas à avancer. La mécanique géométrique permet de décrire comment les contraintes avec l'environnement couplent les articulations au corps. Dans le modèle idéal, une *locomotion naturelle* répète son *état interne* — angles et vitesses articulaires, ainsi que vitesses admissibles du corps — tandis que le corps se translate.

Nous identifions dans cette dynamique interne un canal oscillatoire effectif, associant inertie et énergie potentielle, et un canal propulsif. Comme l'énergie du pendule rebondit entre formes cinétique et potentielle, l'énergie du locomoteur peut rebondir entre le canal oscillatoire effectif et le canal propulsif. Les contraintes idéales organisent sa redistribution sans en fournir. Notre *principe* énonce que, pour qu'une locomotion soit naturelle, ce rebond doit s'équilibrer sur le cycle, sans gain ni perte moyenne dans chacun des canaux. Nous nommons ce transfert *propulsion-oscillator exchange* (POE). L'oscillateur effectif possède lui-même une énergie cinétique et potentielle.

Pour le locomoteur à deux segments (2SEG), un théorème donne une caractérisation nécessaire et suffisante dans le cadre régulier étudié. Pour tester la répétition, nous observons les passages successifs de l'articulation par un angle fixé dans le même sens : une section de Poincaré. Si les vitesses du corps sont identiques à deux passages, un bilan d'énergie POE nul équivaut à retrouver la même vitesse articulaire. Avec une translation nette non nulle, le mouvement est une locomotion naturelle. Une famille régulière de tels cycles et leur phase forment une surface 2D, la *Natural Locomotion Manifold* (NLM). La vitesse moyenne d'avance peut y repérer localement le cycle.

La méthode part d'une oscillation auxiliaire, rétablit le couplage propulsif et calcule un mouvement périodique dans la dynamique interne complète. La continuation numérique suit les cycles voisins. Avec trois segments et deux articulations, le bilan POE reste nécessaire mais ne suffit plus ; le calcul impose la répétition de toutes les composantes de l'état interne. Les résultats distinguent des coordinations en phase et en opposition de phase. Ils éclairent aussi le choix des paramètres mécaniques, en amont du calcul des démarches d'un robot donné.

Avec des pertes, un actionneur doit entretenir le mouvement. Une continuation numérique relie localement un cycle naturel à une réponse entretenue sans changer la mécanique du robot. Une étude distincte, avec d'autres paramètres mécaniques, organise les réponses forcées locomotrices ; leurs composantes régulières définissent localement la *Resonant Locomotion Manifold* (RLM). Comme pour une balançoire, il s'agit d'apporter l'énergie au bon moment. Une *boucle à verrouillage de phase* (PLL) ajuste la fréquence du couple pour maintenir la quadrature : la composante fondamentale de l'angle retarde d'un quart de période sur celle du couple, un repère local de résonance. Une boucle plus lente ajuste l'amplitude pour régler la vitesse moyenne. Les simulations montrent cet entretien sans trajectoire articulaire préenregistrée : la dynamique du robot continue à déterminer le mouvement.

Mots-clés. locomotion naturelle ; modes non linéaires ; mécanique géométrique ; échange d'énergie ; conception mécanique ; résonance ; boucle à verrouillage de phase.

Abstract

Why does running feel more natural than walking ever faster? This intuition invites a robotic question: which gaits does a mechanism favour? This thesis seeks to define, compute and maintain such motions. It first considers an ideal world without dissipation or actuation, separating the organization of motion from the energy needed to sustain it.

Natural oscillations provide a starting point. Near an equilibrium, a linear mode keeps its frequency and coordination as energy increases. At finite amplitude, even the simple pendulum changes frequency; with several joints, coordination can change too. Nonlinear modes describe these families of oscillations: the linear modal plane becomes a curved state-space surface. On a regular modal surface, the action–angle picture gives two intuitive landmarks: action selects the cycle and angle locates motion along it.

Oscillation alone does not produce locomotion. Geometric mechanics describes how environmental constraints couple joint motion to the body. In the ideal model, *natural locomotion* repeats its *internal state*—joint angles and rates, together with admissible body velocities—while the body accumulates a net translation.

Within this internal dynamics, we identify an effective oscillatory channel, combining inertia and potential energy, and a propulsive channel. Just as pendulum energy moves back and forth between kinetic and potential forms, locomotor energy can move between the effective oscillatory and propulsive channels. Ideal constraints organize this redistribution without supplying energy. Our *principle* states that natural locomotion requires this back-and-forth exchange to balance over a cycle, with no mean gain or loss in either channel. We call the transfer *propulsion-oscillator exchange* (POE). The effective oscillator itself contains kinetic and potential energy.

For the two-segment locomotor (2SEG), a theorem establishes a necessary and sufficient characterization in the regular setting studied. To test repetition, we observe successive crossings of a fixed joint angle in the same direction: a Poincaré section. If the body velocities agree at two crossings, zero net POE energy is equivalent to recovery of the same joint rate. With nonzero net translation, the motion is natural locomotion. A regular family of such cycles and their phases forms a two-dimensional surface, the *Natural Locomotion Manifold* (NLM). Mean forward speed can locally identify the cycle on this surface.

The method starts from an auxiliary oscillation, restores propulsive coupling and computes a periodic motion of the full internal dynamics. Numerical continuation follows neighbouring cycles. With three segments and two joints, POE balance remains necessary but is no longer sufficient; the computation enforces repetition of every internal-state component. The results distinguish in-phase and anti-phase organizations. They also inform mechanical-parameter choices, upstream of computing the gaits of a given robot.

With losses, an actuator must sustain the motion. Numerical continuation locally connects one natural cycle to a maintained response without changing the robot's mechanics. A separate study, with different mechanical parameters, organizes forced locomotor responses; their regular components locally define the *Resonant Locomotion Manifold* (RLM). As with a swing, energy must arrive at the right time. A *phase-locked loop* (PLL) adjusts torque frequency to maintain quadrature: the angle fundamental lags the torque fundamental by a quarter-period, a local resonance marker. A slower loop adjusts amplitude to regulate mean speed. Simulations demonstrate this maintenance without a prerecorded joint trajectory: the robot's dynamics continues to determine the motion.

Keywords. natural locomotion; nonlinear modes; geometric mechanics; energy exchange; mechanical design; resonance; phase-locked loop.

Contents

Résumé i

Abstract ii

Notation and conceptual vocabulary vi

I From natural oscillation to natural locomotion 1

1 Letting mechanics organize locomotion 2
- 1.1 A small mechanism, a repeatable movement, a moving body . . . 2
- 1.2 From natural oscillations to an energy-exchange principle . . . 4
- 1.3 A repertoire of gaits and two mechanical-design questions . . . 5
- 1.4 What must the actuator supply? . . . 5
- 1.5 Contributions and their scientific setting . . . 6
- 1.6 How the argument develops . . . 7

2 From natural oscillation to natural locomotion 8
- 2.1 The linear picture: one choreography at every amplitude . . . 8
- 2.2 The pendulum: a nonlinear family with two coordinates . . . 9
- 2.3 Nonlinear modes: coherent motion beyond the linear plane . . . 11
- 2.4 What locomotion adds: a body, an environment and internal dynamics . . . 13
- 2.5 One complete cycle, then a repertoire of cycles . . . 16
- 2.6 From an oscillator's energy cycle to a locomotor's energy cycle . . . 18

II An energy principle and natural gait families 19

3 An energy-exchange principle for natural locomotion 20
- 3.1 Two motions within one internal state . . . 20
- 3.2 Extracting the effective oscillator from the mechanics . . . 21
- 3.3 An auxiliary natural oscillation and its action . . . 23
- 3.4 Power redirected between the two channels . . . 24
- 3.5 Why one effective coordinate gives an exact criterion . . . 25
- 3.6 Where the scalar argument stops . . . 27

4 Computing a repertoire of natural gaits 28
- 4.1 Mechanical model . . . 28
- 4.2 From an organized oscillation to a passive cycle . . . 31
- 4.3 A periodicity problem with independent conditions . . . 32
- 4.4 The computed two-segment repertoire . . . 34

4.5 Changing the mechanics changes the available oscillations . . . 36

5 Two joints, several natural coordinations 39
5.1 What adding an articulation changes . . . 39
5.2 An effective oscillator with two coordinates . . . 40
5.3 Why the scalar energy condition is no longer sufficient . . . 41
5.4 Computing complete internal cycles . . . 42
5.5 Two ways for the same robot to advance . . . 43
5.6 Using the computed motions to guide mechanical design . . . 46
5.7 What the second joint teaches us . . . 48

III Dissipation, resonantly maintained locomotion, and local regulation 50

6 Maintaining locomotion by timed energy input 51
6.1 Replacing lost energy without prescribing a waveform . . . 51
6.2 Following a natural cycle as loss and actuation grow . . . 52
6.3 A broader repertoire of maintained responses . . . 54
6.4 From repeated responses to a locomotion manifold . . . 55
6.5 How excitation changes the locomotor response . . . 57
6.6 Supplying energy at the right phase . . . 57
6.7 A repeated internal motion can keep advancing . . . 59

7 Timing the actuation and regulating mean advance 61
7.1 Being synchronized is not yet choosing a speed . . . 61
7.2 Measuring the phase lag and closing the PLL . . . 62
7.3 Changing mean speed while maintaining phase lock . . . 64
7.4 Closed-loop progression across the rising response segment . . . 66
7.5 What the finite-harmonic model tells us . . . 67
7.6 Accompanying the response: result and scope . . . 67

IV Synthesis, limits, and outlook 69

8 Synthesis: accompanying the mechanics 70
8.1 What the question has become . . . 70
8.2 A principle with a precise scalar realization . . . 70
8.3 A gait repertoire and a mechanical-design question . . . 71
8.4 What the actuator should supply . . . 72
8.5 The next questions opened by the results . . . 72

A A common reduction for continuously constrained locomotors 74
A.1 Admissible velocities and virtual displacements . . . 74
A.2 Projected Lagrange–d'Alembert equation . . . 74
A.3 Why ideal constraints redistribute but do not energize . . . 75
A.4 The two model instances . . . 75
A.5 Scope of the shared method . . . 76

B Two-segment model and conventions 77
B.1 Geometry, velocities, and ideal constraint . . . 77
B.2 Dimensional reduced equations . . . 78

B.3 Dimensionless dictionary and energy check . . . 78
B.4 Linearization for the inertial design relation . . . 79
B.5 Conventions for the numerical results . . . 79

C Scalar coordinates and reproducible periodicity checks 80
C.1 Coefficient functions and opened field . . . 80
C.2 Exchange identity in the implemented convention . . . 81
C.3 The scalar solve and its independent checks . . . 82
C.4 A predecessor-connected 49-node family . . . 83
C.5 The inertia study uses the same physical tests . . . 84

D Three-segment mechanics and periodic-cycle protocol 85
D.1 Exact mechanical model . . . 85
D.2 Modal references and conservative continuation . . . 87
D.3 Independent checks on the seventeen primitive cycles . . . 87
D.4 The separate atlas and local design comparisons . . . 88

E Forced-response equations and numerical verification 90
E.1 Numerical construction of the mechanics-preserving bridge . . . 90
E.2 Equations of the parameter-distinct forced model . . . 91
E.3 Phase extracted from the continued cycles . . . 93
E.4 Peak-speed and quadrature extraction . . . 93
E.5 Population-wide pose reconstruction . . . 94
E.6 Nonlinear shift from the small-amplitude period . . . 95

F Phase–speed controller and sampled finite-harmonic evidence 96
F.1 Rolling phase and speed estimates . . . 96
F.2 Sampled phase and amplitude laws . . . 97
F.3 Local sign tests and the response segment . . . 97
F.4 Stage-transition logic and terminal metrics . . . 98
F.5 Comparing complete cycles and refining the integration step . . . 99
F.6 Closed finite-harmonic controller matrices . . . 100
F.7 What stronger convergence guarantees would require . . . 101

G Local geometry of continued locomotion families 102
G.1 Removing the arbitrary time shift . . . 102
G.2 From a phase-fixed branch to an invariant family . . . 102
G.3 When mean speed is a cycle coordinate . . . 103
G.4 A manifold and its observable image are not the same object . . . 103
G.5 Natural and resonant locomotion objects . . . 104
G.6 What a numerical regularity check establishes . . . 104

Bibliography 105

Notation and conceptual vocabulary

These tables are a consultation aid, not prerequisites for reading. The text introduces the physical ideas before developing their equations. Coefficients specific to a mechanical model are defined in its chapter.

Symbol	Meaning
$q = (g, \boldsymbol{\delta})$	Configuration: rigid pose and internal shape.
$(q, \dot{q})$	Full mechanical state: configuration and velocity.
$g \in \mathrm{SE}(2)$, e_G	Planar pose (position and orientation); identity pose.
δ, $\boldsymbol{\delta}$	One relative joint angle; vector of shape coordinates.
$z = (\boldsymbol{\delta}, w, c)$	Internal state without pose: shape, shape rates and admissible body velocities. Component order is stated for each model.
w, c	Shape-rate vector and independent body-velocity vector.
v, ω, σ	Body-forward speed, yaw rate and one joint rate.
$\bar{v}$	Time average of body-forward speed over a cycle, not generally world-axis displacement divided by period.
T, φ	Period and cycle phase.
J	Action: phase-plane area divided by 2π for a regular one-coordinate conservative oscillator.
Δg	Pose increment reconstructed over a complete internal cycle.
Σ^{χ}	Poincaré section with crossing direction $\chi = \pm 1$.
K, V, E	Kinetic, potential and total mechanical energy.
$M_{\perp}, E_{\perp}$	Effective oscillatory inertia and energy; $E_{\perp}$ includes kinetic and potential energy.
E_{prop}	Complementary propulsive-channel energy.
$P_{\mathrm{POE}}, I_{\mathrm{POE}}$	Propulsion–oscillator exchange power and its time integral; $\dot{E_{\perp}} = P_{\mathrm{POE}}$ in the conservative model.
$A_{\mathrm{exc}}, T_{\mathrm{exc}}, \Omega$	Torque amplitude, excitation period and angular frequency.
$\phi_{\delta\tau}$	Joint-angle fundamental phase minus torque fundamental phase; quadrature target $-\pi/2$.

Concept	Meaning in this thesis
2SEG / 3SEG	Two-/three-segment locomotor, with one/two joints.
IP / AP	In-phase/anti-phase organization: joints move predominantly together/in opposition; not necessarily proportional sinusoids.
Effective oscillator	Auxiliary conservative dynamics derived from the effective inertia and potential; not the full internal dynamics.
POE	Propulsion–oscillator exchange between the effective oscillatory and propulsive channels, both within the internal mechanics.
Pose reconstruction	Integration of position and orientation from body velocity after the internal motion is known.
NLM	Natural Locomotion Manifold: a regular local family of natural cycles swept through phase, two-dimensional in internal-state space.
RLM	Resonant Locomotion Manifold: a regular local transported synchronized-response set; the study uses two-dimensional sheets in orbit–forcing space at fixed forcing phase.
PLL	Phase-locked loop: frequency feedback on measured phase lag.

A cycle repeats the specified internal state, not merely the shape. Locomotion also requires a nontrivial pose increment, with translation for the propulsive examples. Local manifold geometry, orbital stability and attraction are separate questions.

Part I

From natural oscillation to natural locomotion

Chapter 1

Letting mechanics organize locomotion

Why does running feel more natural than walking ever faster? The question suggests that increasing speed need not mean playing the same movement more quickly. A different coordination may become preferable. For a robot, this observation raises a mechanical question before a control law is chosen: *which coordinated movements can the mechanism itself support?* The answer depends on how its masses move, where potential energy is stored, and how interaction with the environment couples its articulations to the motion of the body. Elastic deformation and gravity are familiar sources of potential energy; neither is the only possibility.

This thesis develops a way to identify such movements, compute families of them, and maintain related responses when energy is dissipated. Its starting point is an idealization. We first remove losses and sustained actuation, so that energy supplied at initialization is conserved. This separates two questions that are entangled in a real robot: what organizes a gait, and what energy must be supplied to keep it going? Dissipation and actuation are then reintroduced to study resonant maintenance and speed regulation. The ideal model is a reference for understanding the motion, not a proposal for a perpetually moving physical machine.

1.1 A small mechanism, a repeatable movement, a moving body

The two-segment locomotor in panel (a) of Figure 1.1 will be our main example. An elastic articulation joins a principal body segment to a second segment. A constraint allows the principal segment to move forward and turn but forbids sideways velocity at its reference point. The joint motion and the admissible body velocities are mechanically coupled. Once a motion has been initialized in the conservative model, no actuator tells the joint where to go next: the constrained equations determine its evolution.

Imagine taking a photograph and recording the velocities every time the articulation completes an oscillation. The robot need not be in the same place in successive photographs. Yet it can have the same joint angle, the same joint angular velocity, and the same forward and turning velocities measured in its own frame. The photograph gives the configuration; the velocity measurements complete the mechanical event. The body pose—its position and orientation in the world—has changed. Panel (c) separates these two observations: the internal-state loop closes, while the two placements of the mechanism are distinct. Adding a joint, as in panel (b), will later test how much this organization can change.

We call the repeating collection of joint and admissible body variables the *internal state*, denoted by z. Here "internal" means that absolute placement has been removed; it does not mean that body velocities have been discarded. For the two-segment mechanism,

$$z = (\delta, v, \omega, \sigma), \qquad \sigma = \dot{\delta}, \tag{1.1}$$

where δ is the joint angle, v the forward body velocity, ω the body yaw rate and σ the joint angular

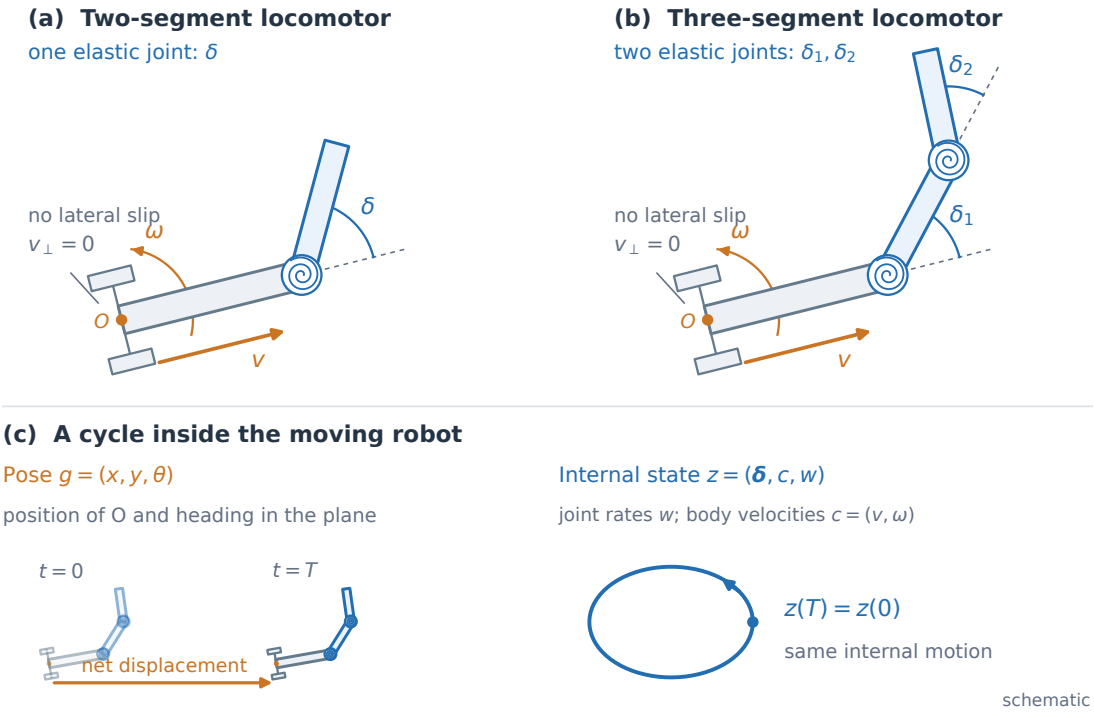


Figure 1.1: Articulation, body motion and placement in the world. The two-segment mechanism has one joint; adding a second joint gives the three-segment example studied later. The internal state includes joint angles and rates *and* the body velocities admitted by the constraints. The body pose is distinct from that state. These are explanatory schematics, not measured robot trajectories.

velocity. The principal segment is also called the *carrier* where this helps describe the model. It is part of the robot, not an external support.

Figure 1.2 turns that schematic into a computed motion. Read panel (a) around one joint angle–rate loop and panel (b) along one reconstructed path. The end of the loop is its beginning; the end of the path is elsewhere. A two-dimensional projection cannot display every component of z; the calculation also checks that the two body velocities repeat. The phase portrait and the spatial view are thus two complementary pictures of one motion: repetition inside the moving robot and displacement in the world.

The formal definition records these physical requirements. Write the configuration as pose and shape,

$$q = (g, \boldsymbol{\delta}), \qquad g \in \mathrm{SE}(2), \tag{1.2}$$

where $\mathrm{SE}(2)$ denotes planar rigid translations and rotations, and $\boldsymbol{\delta}$ collects joint angles. Pose symmetry lets us write the conservative internal dynamics and the reconstruction of pose separately:

$$\dot{z} = f_{\mathrm{cons}}(z), \qquad \dot{g} = g\,\widehat{\Xi(z)}. \tag{1.3}$$

The first equation determines the motion; the second integrates the body velocity $\Xi(z)$ to locate the robot in the world. The hat is the matrix representation of that velocity, explained in Chapter 2.

A natural locomotion cycle is a nonconstant solution of the conservative internal dynamics for which

$$z(T) = z(0), \qquad \Delta g := g(0)^{-1} g(T) \neq e_G, \tag{1.4}$$

with period $T > 0$ and identity pose e_G. The propulsive gaits developed here also have nonzero net translation. This motion is *relative-periodic*: it repeats after allowing for the body's rigid displacement.

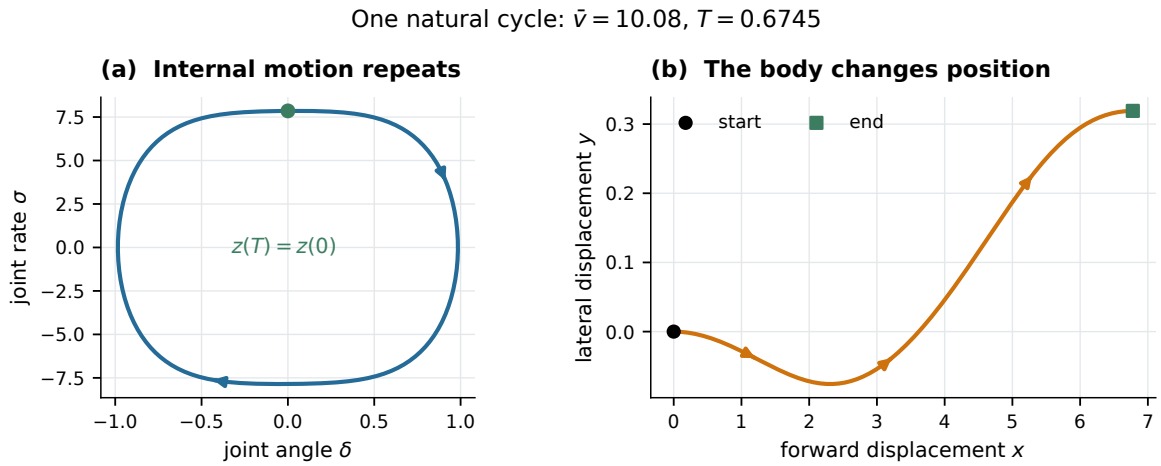


Figure 1.2: One computed natural cycle of the two-segment locomotor. The joint angle–rate projection forms a loop. Integrating the associated body velocity and heading gives a net displacement in the plane. The cycle is selected from the connected family examined in Chapter 4; periodicity is checked in the full internal dynamics, not inferred solely from the projected loop.

Straight motion with a constant internal state is an important reference, but it is not an oscillatory gait merely because an arbitrary period is assigned to it.

The definition identifies a mechanically possible motion. It does not select the most efficient or most stable one. Those are further questions, which require losses, actuation, perturbations or an explicit performance objective.

1.2 From natural oscillations to an energy-exchange principle

A pendulum provides the first intuition. Near its downward equilibrium, the familiar linear approximation predicts a sinusoid with one natural frequency. The exact pendulum is already nonlinear: its period and waveform change with amplitude. Several coupled oscillatory coordinates introduce another freedom, because their relative amplitudes and timing can change together. Nonlinear modal analysis organizes these coherent movements into families rather than treating one small-amplitude waveform as universal.

There is a geometric way to see the family. A single periodic motion is a loop in state space. Neighbouring loops can form a surface. One coordinate selects a loop and another identifies a phase along it. In a regular one-degree-of-freedom conservative oscillator, action–angle coordinates make this description precise: action measures the enclosed phase-space area, and the angle advances as a clock. For several degrees of freedom, a selected nonlinear modal family provides a corresponding local organization.

Locomotion adds the environmental coupling and the moving body. Geometric mechanics supplies tools to describe the permitted velocity directions, remove pose symmetry, derive the internal dynamics and recover displacement. These tools act throughout the construction, not only at its final drawing stage. They also make it possible to examine how the coupled kinetic energy can be reorganized into two meaningful storages.

Within the internal dynamics we identify an *effective oscillatory channel* and a *propulsive channel*. The former combines the potential energy with the inertia effectively involved in joint motion. The latter contains the complementary kinetic storage associated with propulsion. Because the body and joints are coupled, this is not a separation of robot parts into two independent machines. Nor is it simply another name for kinetic and potential energy: the effective oscillator itself has both.

> Just as a pendulum exchanges energy between its kinetic and potential forms, a locomotor can exchange energy between its effective oscillatory and propulsive channels. For a natural locomotion cycle, this exchange must balance: neither channel has a mean energy gain or loss over the cycle.

This is the *propulsion–oscillator exchange* (POE) principle developed in this thesis. Ideal environmental constraints organize the redistribution without supplying total energy. The two channel energies can change during the motion even though their sum stays constant. Their recovery after a cycle therefore expresses more than conservation of that sum. The construction in Chapter 3 defines the storages and POE power from the mechanical equations before a periodic trajectory is known.

For the two-segment locomotor, this viewpoint yields an exact result. Observe the joint on a Poincaré section: each time it crosses a chosen angle in the same direction. Suppose the body velocities agree at two successive crossings. The theorem shows that zero net POE energy between them is equivalent to recovery of the same joint velocity. Together these conditions make the internal state repeat. With the displacement requirement, they characterize the natural cycles in the regular setting considered. The proof, its hypotheses and its geometric meaning are all developed in the main text.

1.3 A repertoire of gaits and two mechanical-design questions

One cycle answers how the robot can move. A family answers how its motion can vary. A regular one-parameter family of natural cycles, swept through their phases, forms a local two-dimensional surface in internal state space: the *Natural Locomotion Manifold* (NLM). The word manifold describes that underlying organization, not every possible trajectory or the appearance of a smooth plot.

Mean forward speed is an especially useful physical label where it varies regularly along the family. Then a reader can picture a repertoire in which one choice determines the mean advance and phase locates the robot on the chosen gait. The mean is taken from body-frame forward velocity over the cycle; it need not equal world-frame displacement divided by the period when the robot turns. Under passive evolution, the robot follows its chosen cycle. Moving between cycles is a separate actuation or initialization problem, rather than an automatic consequence of the surface geometry.

The calculation addresses two related robotic questions. For a given robot, it seeks the natural gaits allowed by that robot's mechanics. Before the design is chosen, the same construction can help investigate which masses, inertias, lengths or elastic properties permit useful gaits. These are two different variations: changing a gait without changing the mechanism, and changing the mechanism to alter its repertoire. The two-segment inertia analysis and the three-segment design comparisons develop both uses.

With two joints, coordinated movements can be predominantly in phase or in opposition of phase. Recovering the same effective-oscillator energy then no longer fixes every joint coordinate and velocity. The method retains the complete periodicity conditions while using modal organization and POE to guide and interpret the solutions. This is a constructive extension of the scalar result, not a claim that a single energy equation controls an arbitrarily large mechanism.

1.4 What must the actuator supply?

The ideal reference separates the organization of motion from its energetic maintenance. In the presence of losses, that separation becomes useful: the actuator must replace dissipated energy, but it need not prescribe every detail of the joint waveform. Pushing a swing at the right time illustrates the distinction. The push supplies energy and timing; it does not specify the swing's position at every instant.

A numerical continuation first follows one conservative locomotor cycle as dissipation and periodic actuation are introduced, while retaining that robot's inertia, geometry and potential. The resulting waveform and period change. This gives a local example of a natural motion becoming a maintained response. A separate forced-response study explores a wider family on a related mechanism with different parameters. Its regular locomotor response components define local charts of a *Resonant Locomotion Manifold* (RLM). These two studies answer complementary questions and are kept distinct.

The final step uses two measured quantities to regulate motion. A *phase-locked loop* (PLL) adjusts torque frequency according to the phase lag of the joint-angle fundamental relative to the torque. A slower loop adjusts torque amplitude according to mean forward speed. Quadrature, where the angle fundamental lags by a quarter-period, provides a useful local timing target. The complete nonlinear waveform remains the robot's response, not a prerecorded trajectory sent to its joint.

1.5 Contributions and their scientific setting

The contributions form one argument, with distinct analytical and numerical parts:

1. A mechanically defined **POE principle**: the effective oscillatory and propulsive storages make the energy exchanged during natural locomotion physically and mathematically explicit.
2. An **exact 2SEG characterization**: under the stated oriented- section conditions, zero net POE energy is equivalent to the remaining joint-velocity periodicity condition. This characterizes the natural cycles of the corresponding local NLM when the other conditions hold.
3. A **method for computing natural gait families**, with a connected two-segment numerical family, a demonstrated scalar/vector boundary and two three-segment joint organizations at one shared design. The design studies explain how the available motions inform mechanical choices.
4. A **numerical passage to maintained locomotion**: one local conservative-to-forced continuation and a separate nonlinear locomotor response study, with a precise local RLM interpretation.
5. **Local phase–speed regulation**: PLL-based timing and slower amplitude adjustment maintain simulated locomotion without joint-waveform prescription.

1.5.1 The established ideas that make these results possible

Passive-dynamic walking and its actuated descendants show that robot mechanics can organize much of a gait (McGeer, 1990; Collins et al., 2005). In a downhill walker, gravitational input can balance impact losses; "passive" does not by itself mean conservative. Conservative legged models also reveal families and connections between familiar gait types (Remy, 2011; Gan et al., 2018; Raff et al., 2022). These results make a repertoire-based question meaningful before any particular locomotor is chosen.

Nonlinear modal analysis provides the finite-amplitude oscillatory organization used here. Its developments in nonlinear dynamics and structural dynamics explain changing frequencies, waveforms and modal surfaces (Rosenberg, 1966; Shaw and Pierre, 1993; Kerschen et al., 2009; Peeters et al., 2009). Robotics has extended this viewpoint to the excitation and stabilization of mechanically organized motions, including experimental nonlinear-mode excitation and natural-motion-family control (Albu-Schäffer and Della Santina, 2020; Della Santina and Albu-Schäffer, 2021; Bjelonic et al., 2022; Calzolari et al., 2025). The present thesis uses that language to organize oscillation while explicitly retaining the body-velocity dynamics needed for locomotion. In particular, Calzolari et al. (2025) already compute and excite families of passive quadruped gaits. Neither the existence of a gait family nor its excitation is claimed here as a new idea. The specific addition is the effective-oscillator/propulsion energy partition, its scalar section equivalence, and its use in the constrained locomotors developed below.

Geometric and nonholonomic mechanics describe the interaction of shape, symmetry, momentum and environmental constraints (Kelly and Murray, 1995; Ostrowski and Burdick, 1998; Bloch, 2015; Hatton and Choset, 2011). Reduced locomotion models with passive internal degrees of freedom and internal rotors already bring these ingredients together (Boyer and Belkhiri, 2014; Fedonyuk and Tallapragada, 2019). Relative-periodic-orbit theory makes repetition with displacement precise (Fassò et al., 2020; Eldering and Jacobs, 2016). The contribution here is the specific energetic construction, exact scalar criterion and resulting locomotor calculations within this established mechanical framework, not the invention of reduction or reconstruction.

Shooting and continuation provide mature means to compute periodic solutions (Peeters et al., 2009; Kuznetsov, 2004). Forced nonlinear response, backbone theory and phase-resonance testing explain how free oscillatory organization can inform maintained responses (Cenedese and Haller, 2020a; Peeters et al., 2011; Peter and Leine, 2017; Denis et al., 2018). Control-based continuation is another established way to explore nonlinear responses (Renson et al., 2016; Barton, 2017). The thesis develops a locomotor use of these tools and separates its offline family computation from its online phase–speed feedback.

Nor is mechanics-aware control confined to full-trajectory tracking. Orbital stabilization and hybrid zero dynamics provide important alternatives (Canudas-de Wit et al., 2002; Westervelt et al., 2003). Here the particular choice is to regulate the timing and energetic supply of a locomotor response. Biological and resonant-robot studies motivate that choice without making resonance a universal explanation of efficiency (Tytell et al., 2014; Ramananarivo et al., 2011; Schönebaum et al., 2021).

Parts of this research have appeared in co-authored publications on the two-segment mechanism and forced locomotion (Rajaomarosata et al., 2025a,b), and in an earlier preprint of the principle-and-method study (Mortel et al., 2026). The preprint belongs to the development of the same work; it is not independent validation of the results presented here. The dissertation develops the reasoning and connects the conservative, design and maintained-motion contributions in one account.

1.6 How the argument develops

Chapter 2 follows the concepts from a familiar oscillator to a moving articulated body. It introduces action, phase, modal geometry, admissible velocities and relative periodicity where they clarify that transition. Chapter 3 then constructs the two energy channels and proves the two-segment result. Only after the principle and its meaning are established does Chapter 4 develop the computation, the numerical gait repertoire and the two-segment mechanical-design analysis.

Chapter 5 adds the second joint, explains why the scalar condition no longer suffices, and examines the two joint organizations and their design dependence. Chapter 6 reintroduces energy loss and timed input. Chapter 7 develops the PLL and mean-speed regulator, and the final chapter draws together what the mechanics selects and what the actuator supplies.

The results concern declared continuous constrained models and are analytical or numerical. Impacts, contact switching and physical hardware validation require additional modelling or experiments. The main text contains the reasoning needed to understand the contributions; the appendices retain coefficient derivations, numerical protocols and local geometry for readers who want to reproduce or examine them more closely.

Chapter 2

From natural oscillation to natural locomotion

A pendulum and a locomotor both coordinate motion through mechanics, but only one changes its place in the world. To connect them, we need two complementary ideas. Nonlinear vibration theory explains how a free oscillation can retain an organized rhythm while its frequency and choreography change with energy. Geometric mechanics explains how internal deformation and environmental constraints participate in the motion of a body. Together they suggest how to look for a natural locomotor rhythm without prescribing its waveform.

We begin with the familiar linear mode, enlarge the picture to finite amplitudes, and then add the ingredients needed for locomotion. The aim is to make the construction in the next chapter intelligible: what is periodic, which energy is being exchanged, and why the one-joint locomotor admits a scalar characterization whereas the two-joint model requires more information. Throughout, the closed internal loop and open pose path of Figure 1.2 remain the locomotor reference.

2.1 The linear picture: one choreography at every amplitude

For a small displacement $\boldsymbol{u}$ from a stable equilibrium, the linear approximation to an undamped mechanical system is

$$M_0\ddot{\boldsymbol{u}} + K_0\boldsymbol{u} = 0, \qquad K_0\boldsymbol{\phi}_j = \omega_j^2 M_0\boldsymbol{\phi}_j. \tag{2.1}$$

The mass matrix M_0 and stiffness matrix K_0 are evaluated at equilibrium. We consider an oscillatory mode with $\omega_j > 0$; a zero-frequency rigid motion would not trace the closed ellipses described below. The eigenvector $\boldsymbol{\phi}_j$ specifies a mode's relative displacements: each coordinate follows the same sinusoid with a fixed ratio and sign. An individual modal motion is

$$\boldsymbol{u}(t) = a\boldsymbol{\phi}_j \cos(\omega_j t + \varphi_0), \qquad \dot{\boldsymbol{u}}(t) = -a\omega_j\boldsymbol{\phi}_j \sin(\omega_j t + \varphi_0). \tag{2.2}$$

Changing a changes the energy, but neither the frequency ω_j nor the coordination encoded by $\boldsymbol{\phi}_j$. In this sense a linear mode has one choreography, enlarged or reduced without being reshaped.

There is an equally useful geometric picture. All pairs $(\boldsymbol{u}, \dot{\boldsymbol{u}})$ associated with this mode lie in the two-dimensional plane spanned by $(\boldsymbol{\phi}_j, 0)$ and $(0, \boldsymbol{\phi}_j)$. An energy level cuts that plane along an ellipse; one orbit travels around the ellipse. One parameter selects the ellipse, and another locates the state along it. The plane is invariant under the linear equations: a state initially in it remains there.

The modal plane is exact for the linear model. Linearization retains the first variation of restoring and inertial forces at equilibrium, giving a useful description of small vibrations in structural dynamics and robotics (Kerschen et al., 2009; Albu-Schäffer and Della Santina, 2020). At finite amplitudes, nonlinear forces can change both the oscillation period and the coordination of the moving parts.

2.2 The pendulum: a nonlinear family with two coordinates

The simple pendulum is already nonlinear. For angular displacement x, its restoring torque is proportional to $\sin x$, not to x. Only its small-angle approximation has a fixed natural frequency. A large libration still repeats without actuation in the ideal conservative model, but it takes longer than a small one. Natural oscillation therefore does not mean a sinusoid at an immutable frequency.

2.2.1 Energy selects a closed loop

For one coordinate x, let $M(x) > 0$ be the inertia, $U(x)$ the potential, and $p = M(x)\dot{x}$ the conjugate momentum. In these coordinates the Hamiltonian is the mechanical energy:

$$H(x, p) = \frac{p^2}{2M(x)} + U(x) = E. \tag{2.3}$$

In a regular potential well, a bounded energy level has two turning points $x_-(E)$ and $x_+(E)$. Between them,

$$p_\pm(x; E) = \pm\sqrt{2M(x)\{E - U(x)\}}. \tag{2.4}$$

The two signs describe opposite directions of motion. They join at the turning points and form a closed curve in the position–momentum plane. One circuit is one natural oscillation. Potential energy increases as the pendulum climbs; kinetic energy increases as it falls. The energy repeatedly passes between these two forms, with no net gain or loss over a cycle. In Figure 2.1, follow one energy level from panel (a) to panel (b): its two turning points become the ends of the positive- and negative-momentum branches. Changing the energy selects a different complete loop, not a different point on the same motion.

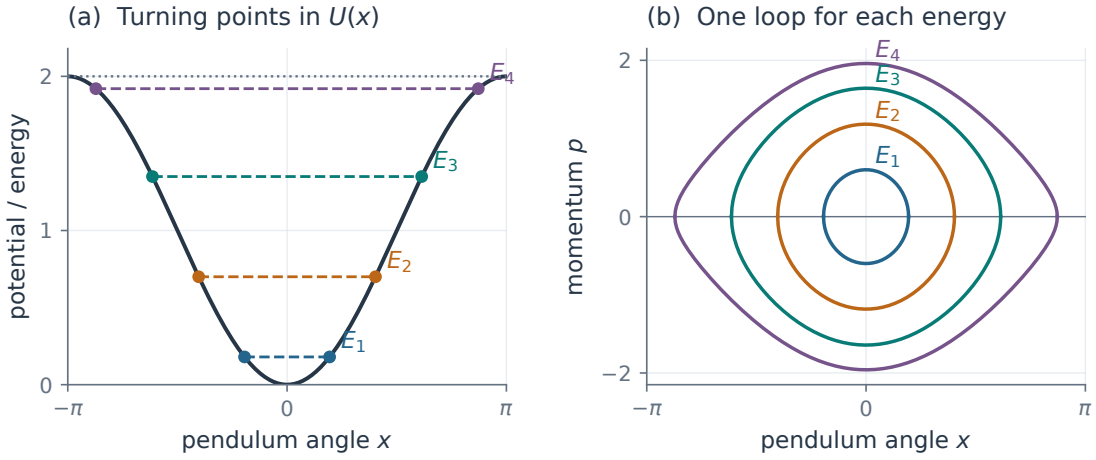


Figure 2.1: Energy levels organize a family of pendulum librations. The potential determines the two turning points (left); the positive and negative momentum branches complete a closed state-space orbit (right). Matching labels identify the same energy in both panels. This analytical pendulum illustration uses $I = \omega_0 = 1$, so $U(x) = 1 - \cos x$.

The time required to traverse the orbit follows from $\dot{x} = p/M(x)$:

$$T(E) = 2\int_{x_-(E)}^{x_+(E)} \sqrt{\frac{M(x)}{2\{E - U(x)\}}}\,\mathrm{d}x, \qquad \Omega(E) = \frac{2\pi}{T(E)}. \tag{2.5}$$

Near a turning point the speed is small, so a short displacement takes a relatively long time. This also happens in a harmonic oscillation; it is not by itself a signature of nonlinearity. What changes in

general is the relation between position and speed over the whole orbit, through $U(x)$ and $M(x)$. Its waveform need not be sinusoidal and its period can depend on energy. Constant inertia and a quadratic potential recover the familiar energy-independent period of the harmonic case.

For the simple pendulum specifically, x is its angular displacement and

$$H(x,p) = \frac{p^2}{2I} + I\omega_0^2(1-\cos x), \qquad \omega_0 = \sqrt{g/l}. \tag{2.6}$$

Here I is the inertia about the pivot, l the pendulum length and g gravitational acceleration. For maximum angle $0 < a < \pi$,

$$T(a) = \frac{4}{\omega_0}\,\mathrm{K}\big(\sin(a/2)\big), \qquad \mathrm{K}(k) = \int_0^{\pi/2} \frac{\mathrm{d}\psi}{\sqrt{1-k^2\sin^2\psi}}. \tag{2.7}$$

The period tends to $2\pi/\omega_0$ at small amplitude and diverges on approaching the upright separatrix. The family thus contains a continuum of natural frequencies, not a single frequency with different amplitudes.

2.2.2 Action chooses the orbit; angle tells the time within it

Energy labels neighbouring closed loops. A closely related label, useful in Hamiltonian mechanics, is the *action*

$$J(E) = \frac{1}{2\pi}\oint_{H=E} p\,\mathrm{d}x. \tag{2.8}$$

It is the area enclosed by the loop divided by 2π, with the orientation of the mechanical flow. The shaded region in Figure 2.2(a) makes this definition visible: action measures a complete state-space oscillation, not just its maximum displacement.

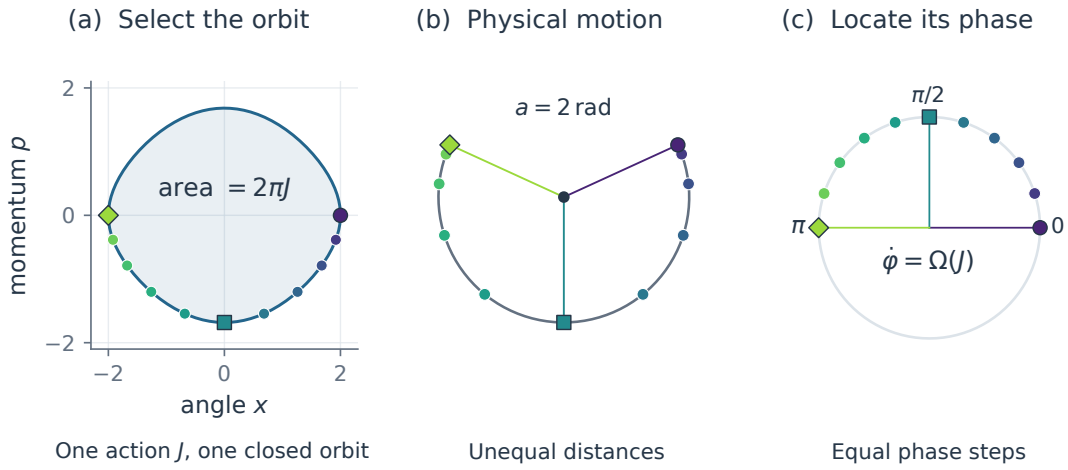


Figure 2.2: One orbit, two coordinates: action and phase. The shaded phase-plane area is $2\pi J$ (a); physical position (b) and uniform phase (c) describe the same motion. Eleven points mark equally spaced instants over one half-period of the exact pendulum libration with amplitude 2 rad and $I = \omega_0 = 1$. Circle, square and diamond identify its start, middle and end in all three panels. Phase is a clock, not a new physical articulation.

To connect this area to the clock, write $J = \pi^{-1} \int_{x_-}^{x_+} p_+ \, \mathrm{d}x$ and note that $\partial p_+ / \partial E = M/p_+$. At simple turning points, the endpoint terms vanish because $p_+ = 0$, and the remaining integral is convergent. Differentiating thus gives

$$\frac{\mathrm{d}J}{\mathrm{d}E} = \frac{1}{\pi} \int_{x_-}^{x_+} \frac{M(x)}{p_+(x;E)} \, \mathrm{d}x = \frac{T(E)}{2\pi} > 0. \tag{2.9}$$

Thus action and energy label the same local family. An angle $\varphi \in S^1$ can be chosen to advance uniformly around each orbit:

$$\dot{J} = 0, \qquad \dot{\varphi} = \frac{\mathrm{d}E}{\mathrm{d}J} = \Omega(J) = \frac{2\pi}{T(J)}. \tag{2.10}$$

These are *action–angle coordinates* (Arnold, 1989). The action chooses which oscillation occurs; the angle locates the state within it. Uniform angle does not imply uniform physical speed:

$$x(t) = X\big(J_0, \varphi_0 + \Omega(J_0)t\big), \tag{2.11}$$

where X is periodic in its second argument but need not be a cosine. The corresponding points in panels (a)–(c) of Figure 2.2 all represent the same instants. During the illustrated half-cycle, equal time steps are unevenly spaced along the pendulum arc but evenly spaced around the phase clock. The action chooses the orbit; the phase records progress along that orbit.

The coordinates are canonical: with the present order convention, $\mathrm{d}x \wedge \mathrm{d}p = \mathrm{d}\varphi \wedge \mathrm{d}J$. This area-preserving property distinguishes them from an arbitrary drawing of the orbit as a circle. Their domain is a regular annulus of bounded periodic motions. Equilibrium and separatrix are boundaries of this construction, not ordinary points of one global chart.

This picture will matter twice. It provides an intuitive description of a two-dimensional family of natural oscillations. It also supplies the one-coordinate energy argument used for the effective oscillator inside 2SEG. Before making that connection, we need to understand what changes when the oscillator itself has several coordinates.

2.3 Nonlinear modes: coherent motion beyond the linear plane

With two freely oscillating coordinates, one energy value no longer selects one orbit. More generally, a regular energy surface in a $2n$-dimensional mechanical state space has dimension $2n - 1$. For the double pendulum this leaves three dimensions, whereas a periodic orbit is only one-dimensional. The same energy may support different periodic motions, quasiperiodic motions and, in appropriate regimes, chaotic trajectories. Energy conservation alone does not organize them into a single choreography.

Nonlinear normal modes provide such an organization for selected motions. This is an established contribution of nonlinear vibration and structural dynamics, with extensions and uses in robotics (Rosenberg, 1966; Shaw and Pierre, 1993; Kerschen et al., 2009; Peeters et al., 2009; Albu-Schäffer and Della Santina, 2020). The conservative periodic-family view is the one most directly useful here: follow a coherent family of periodic motions while allowing the frequency, relative amplitudes and waveform to evolve.

2.3.1 How a modal plane becomes a curved surface

In the linear model, exciting one eigenvector leaves every other linear mode at zero. In the nonlinear equations, the same initial configuration can activate coupling terms that push the motion away from that plane. An amplified linear solution is therefore generally not an exact finite-amplitude solution.

A nonlinear periodic family accounts for those terms. Neighbouring cycles adjust their period and coordination so that they solve the nonlinear equations. Sweeping every cycle through its phase produces a two-dimensional surface: one coordinate chooses the cycle, and the other chooses the point on it. On a regular portion this may be written

$$W_j : (a, \varphi) \longmapsto \big(\boldsymbol{u}, \dot{\boldsymbol{u}}\big), \qquad \varphi \in S^1. \tag{2.12}$$

For a family emerging from a linear mode under the relevant regularity and non-resonance hypotheses, the surface is tangent to that modal plane at equilibrium; away from equilibrium it generally curves. It is invariant when the vector field is tangent to it: a freely moving state on a selected orbit remains on that orbit. In Figure 2.3, first follow one coloured loop, then compare neighbouring loops with the grey linear plane. Their departure from that plane shows why increasing energy need not simply enlarge a linear choreography.

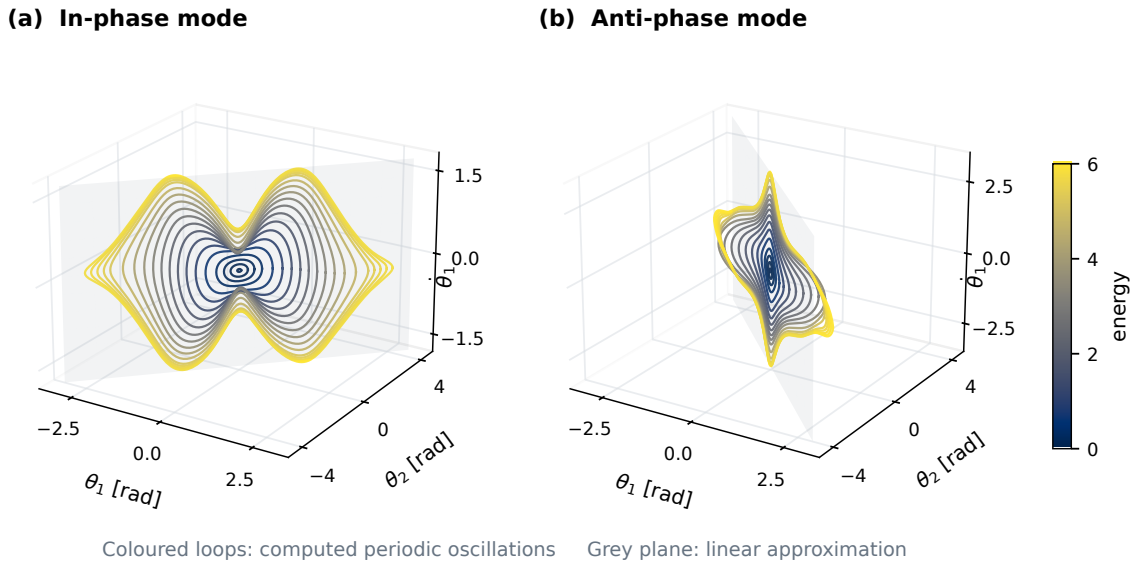


Figure 2.3: From a linear plane to finite-amplitude modal geometry, illustrated by a symmetric double pendulum. The two sectors are organized by in-phase and anti-phase small-angle directions. The curves are stored numerical periodic motions outlining the nonlinear family; the grey surface is the tangent linear plane. No nonlinear trajectories are filled in between the sampled loops. Each panel projects three coordinates of the four-dimensional mechanical state.

For equal point masses and equal link lengths, the two small-angle link directions, expressed in the absolute angles of the two links, are proportional to $(1, \sqrt{2})^\mathsf{T}$ and $(1, -\sqrt{2})^\mathsf{T}$. In the first, the links move predominantly together; in the second they move predominantly in opposition. At finite amplitudes these labels describe organization, not fixed amplitude ratios or perfectly copied sine waves. Comparing the two panels therefore shows two different organizations of the same mechanism, not two drawings of one energy surface.

A regular modal family can be pictured as an island of organized motion within a much larger set of possibilities. The double pendulum makes the contrast particularly vivid because its surrounding dynamics can be complicated. The image does not mean that every surrounding motion is chaotic, nor that a conservative modal surface attracts nearby states. Periodicity, stability and attraction are different properties.

2.3.2 A family label and a phase, without global integrability

A backbone curve records a modal family's frequency $\Omega_j = 2\pi/T_j$ against energy, action or amplitude. Unlike the linear eigenfrequency, this frequency may change along the family. The backbone is a

summary of complete cycles, not a replacement for their waveforms.

On a regular two-dimensional Hamiltonian modal surface for which the restricted symplectic structure is nondegenerate, action–angle coordinates give the same local picture as for the pendulum:

$$\dot{J}_j = 0, \qquad \dot{\varphi}_j = \Omega_j(J_j). \tag{2.13}$$

This describes the selected coherent motions. It does not imply that the whole double pendulum, or an arbitrary many-coordinate mechanism, has global action–angle coordinates. Invariant-surface existence and smoothness require appropriate spectral and regularity hypotheses; the periodic-family and invariant-manifold perspectives must be related with those hypotheses in view (Shaw and Pierre, 1993; Haller and Ponsioen, 2016).

For locomotion, this distinction anticipates an important difference. The auxiliary effective oscillator of 2SEG has one coordinate; regular bounded energy curves organize its oscillations directly. The corresponding oscillator of 3SEG has two coordinates; a modal organization must additionally be selected. Neither auxiliary oscillator yet describes the complete locomotor dynamics. Both must be coupled to the propulsive motion. Geometric mechanics supplies the next ingredients.

2.4 What locomotion adds: a body, an environment and internal dynamics

The vibration examples have a fixed support. Their coordinates describe deformation or motion relative to that support. A locomotor must also move its body through the environment. The articulations and the body cannot be analysed as unrelated mechanisms: their admissible motions are coupled by the interaction with the environment.

Geometric mechanics makes this coupling explicit. Its contributions include configuration symmetry, connections between shape and body motion, momentum-dependent reconstruction, and constrained equations of motion (Kelly and Murray, 1995; Ostrowski and Burdick, 1998; Bloch, 2015; Hatton and Choset, 2011). These tools are not restricted to prescribing a shape loop. Inertial locomotion models also treat passive internal degrees of freedom and the dynamics that selects their motion (Boyer and Belkhiri, 2014; Fedonyuk and Tallapragada, 2019). We use this established framework to identify a particular energetic organization of natural locomotion.

2.4.1 Configuration, full state and internal state

Write the configuration as

$$q = (g, \boldsymbol{\delta}), \qquad g \in G, \qquad \boldsymbol{\delta} \in \mathcal{S}. \tag{2.14}$$

The pose g places and orients the whole body in the world. The shape $\boldsymbol{\delta}$ describes its internal configuration, such as its joint angles. For planar locomotion, $G = \mathrm{SE}(2)$: two translations and one rotation. The full mechanical state contains configuration and velocity, $(q, \dot{q})$. Neither configuration alone nor shape alone determines the subsequent inertial motion.

When the mechanical laws are unchanged by a common rigid displacement, the equations need not retain absolute position and heading. Expressing body velocities in a body-attached frame gives an *internal state*

$$z = (\boldsymbol{\delta}, w, c), \qquad w = \dot{\boldsymbol{\delta}}. \tag{2.15}$$

Here c collects the independent body velocities permitted by the constraints. In the applications below, $c = (v, \omega)$ consists of forward body speed and yaw rate. The absence of pose from z does *not* mean the absence of body motion: the permitted body velocities are part of the internal dynamics. Table 2.1 keeps this removal of absolute placement separate from the energetic partition that will be made within the internal state.

Table 2.1: Four descriptions with different purposes. Removing pose symmetry retains the body velocities needed by the dynamics. The energetic partition introduced next is a different operation within that internal dynamics.

Description	Variables	Question answered
Configuration	$q = (g, \boldsymbol{\delta})$	Where is the robot and how is it shaped?
Full state	$(q, \dot{q})$	What is its configuration and motion in the world?
Internal state	$z = (\boldsymbol{\delta}, w, c)$	What determines the motion independently of absolute pose?
Two energy channels	$E_{\perp}$, E_{prop}	How is energy redistributed within the internal mechanics?

For 2SEG, there are four configuration coordinates, three independent admissible velocity coordinates, and one joint. For 3SEG, these numbers are five, four and two. Thus the scalar oscillator used for 2SEG does not make the entire locomotor a one-degree-of-freedom system.

2.4.2 Constraints redirect energy without supplying it

The general motivation concerns interaction with the environment. For the two continuous mechanical applications, that interaction is represented specifically by ideal homogeneous constraints on velocities,

$$A(q)\dot{q} = 0. \tag{2.16}$$

On a constant-rank domain, a basis $H(q)$ of the allowed velocity directions gives

$$\dot{q} = H(q)\eta, \qquad A(q)H(q) = 0, \qquad \eta = (c, w). \tag{2.17}$$

For instance, a no-sideways-slip constraint forbids one lateral velocity while allowing longitudinal motion and rotation. Its reaction changes which velocities the mechanism can acquire.

With Lagrangian $L = K - V$, the Lagrange–d'Alembert equations are

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = Q + A(q)^{\mathsf{T}}\lambda, \tag{2.18}$$

where Q denotes applied generalized forces. Projecting onto the allowed directions eliminates the unknown ideal reactions:

$$H(q)^{\mathsf{T}}\left(\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q\right) = 0. \tag{2.19}$$

This is the common construction behind both locomotors; the detailed coefficient derivation is in appendix A.

The same geometry establishes the power of the ideal reaction:

$$\dot{q}^{\mathsf{T}}A(q)^{\mathsf{T}}\lambda = \big(A(q)\dot{q}\big)^{\mathsf{T}}\lambda = 0. \tag{2.20}$$

The reaction supplies no total mechanical energy. It nevertheless redirects motion and, through the coupled inertial dynamics, organizes redistribution between the effective oscillatory and propulsive channels derived in the next chapter. Zero external work does not mean that the channels evolve independently.

This calculation assumes a smooth, continuously active, ideal, constant-rank velocity constraint and an autonomous conservative model. Impacts, changing contacts and dissipative environmental forces need their own energy balances. The principle motivates a broader question; the proofs and computations retain the precise scope of the declared models.

2.4.3 Projected momentum connects geometry to the oscillator

Substituting the admissible velocities into the kinetic energy produces a positive-definite reduced mass matrix M_r:

$$K = \frac{1}{2}\begin{bmatrix} c \\ w \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} M_{cc} & M_{cw} \\ M_{cw}^{\mathsf{T}} & M_{ww} \end{bmatrix} \begin{bmatrix} c \\ w \end{bmatrix}, \qquad \begin{bmatrix} \pi_c \\ \pi_w \end{bmatrix} = M_r \begin{bmatrix} c \\ w \end{bmatrix}. \tag{2.21}$$

The momentum here is projected onto the admissible velocity directions. Such components are often called *pseudo-momenta*: the velocity coordinates need not be derivatives of independent configuration coordinates. They should not automatically be treated as canonical momenta of the whole nonholonomic system.

The combination

$$p_\perp = \pi_w - M_{cw}^{\mathsf{T}} M_{cc}^{-1} \pi_c = \underbrace{\left(M_{ww} - M_{cw}^{\mathsf{T}} M_{cc}^{-1} M_{cw}\right)}_{M_\perp} w \tag{2.22}$$

removes the body-momentum contribution from the shape momentum. The positive Schur complement $M_\perp$ is the inertia of the effective oscillator. For its one-coordinate 2SEG auxiliary dynamics, $p_\perp = M_\perp(\delta)\sigma$ is the momentum entering the action integral:

$$J_\perp(E) = \frac{1}{2\pi} \oint p_\perp \, \mathrm{d}\delta. \tag{2.23}$$

This makes the pendulum analogy concrete. The action is built from the effective momentum obtained through the constrained mechanical construction, not from an unrelated geometric label. The auxiliary oscillator has its own canonical position–momentum pair; no global Hamiltonian action–angle structure is asserted for the full locomotor.

2.4.4 From internal motion to displacement

Let ξ denote the body twist, represented by a matrix $\widehat{\xi}$. The relation between body velocity and world pose is

$$\widehat{\xi} = g^{-1}\dot{g}, \qquad \dot{g} = g\widehat{\xi}. \tag{2.24}$$

The mechanical and pose equations can then be organized as

$$\dot{z} = f_0(z), \qquad \xi = \Xi(z), \qquad \dot{g} = g\widehat{\Xi(z)}. \tag{2.25}$$

The first equation determines the timed internal motion; the last integrates where that motion carries the body. This last operation is called *pose reconstruction*.

For a purely kinematic model, a connection can sometimes express $\xi = -\mathcal{A}(\boldsymbol{\delta})\dot{\boldsymbol{\delta}}$. In the inertial models used here, body velocities remain dynamical variables, so reconstruction depends on the internal state, not on a shape loop alone. A useful kinematic example nonetheless shows why a closed internal deformation need not produce a closed pose path. For two shape coordinates, suppose

$$\dot{x} = \frac{1}{2}(\delta_2\dot{\delta}_1 - \delta_1\dot{\delta}_2). \tag{2.26}$$

A counterclockwise closed shape curve enclosing area $\mathcal{A}_S$ then produces

$$\Delta x = \frac{1}{2}\oint (\delta_2 \, \mathrm{d}\delta_1 - \delta_1 \, \mathrm{d}\delta_2) = -\mathcal{A}_S. \tag{2.27}$$

This is a simple geometric phase: deformation repeats while position changes. It describes the displacement of a supplied loop; selecting that loop as a free mechanical motion requires the dynamics as well.

For planar inertial locomotion, heading also changes. With pose $g = (R, p)$, composition is

$$(R_1, p_1)(R_2, p_2) = (R_1 R_2, p_1 + R_1 p_2). \tag{2.28}$$

A translation is rotated by the preceding heading. Consequently the ordered history of body velocities matters; their separate time integrals do not in general determine the world displacement.

2.4.5 What must repeat in natural locomotion

A nonconstant conservative internal motion is periodic if $z(T) = z(0)$ for some $T > 0$. It is a transporting relative-periodic motion when

$$\dot{z} = f_0(z), \qquad z(T) = z(0), \qquad \Delta g = g(0)^{-1} g(T) \neq e_G. \tag{2.29}$$

The equations express three physical requirements: compatibility with the passive mechanics, repetition of the complete internal state, and nontrivial body motion over the cycle. In the propulsive examples, translation is checked as well, distinguishing advance from rotation in place. Relative periodicity is an established geometric notion (Fassò et al., 2020; Eldering and Jacobs, 2016); the energy principle developed here characterizes a specific way these motions are organized.

Shape recurrence alone would be too weak. The same joint angle can be crossed in opposite directions; the same joint motion can coexist with different body velocities. Both change the subsequent mechanical evolution. Conversely, imposing the initial world pose again would exclude the net displacement we seek. A natural cycle repeats its internal state while leaving its pose free to accumulate:

$$g(nT) = g(0)(\Delta g)^n. \tag{2.30}$$

A straight increment produces repeated translation; an increment with rotation can yield a turning path. This is the formal counterpart of the two views in Figure 1.2: a loop in internal-state space supplies one displacement increment in the world.

The resulting cycle resembles an ordinary oscillator in its internal state, but does more than oscillate: it transports the mechanism. The next section provides a precise way to recognize repetition without having to compare every possible starting phase.

2.5 One complete cycle, then a repertoire of cycles

2.5.1 A Poincaré section is a repeatable observation

Imagine taking a photograph whenever a joint passes through its central position in the same direction. For a simple oscillation this event occurs once per complete cycle. The photograph fixes where the clock starts; measuring the other internal variables tells us whether the next passage really repeats the same mechanical state.

A *Poincaré section* formalizes that observation. For an autonomous field $\dot{z} = f(z)$ and a smooth event function h, choose

$$\Sigma^+ = \{z : h(z) = 0,\ \mathrm{D}h(z) f(z) > 0\}. \tag{2.31}$$

The positive derivative selects a crossing direction and excludes tangency. If $\tau(z) > 0$ is the next finite crossing time in that direction, the Poincaré map is

$$P(z) = \Phi^{\tau(z)}(z). \tag{2.32}$$

A fixed point $P(z_\star) = z_\star$ is a periodic internal motion (Kuznetsov, 2004). The section removes the arbitrary time origin; it does not apply a physical constraint or force to the robot. In Figure 2.4, both marked points lie on the geometric section, but only the lower crossing has the selected direction.

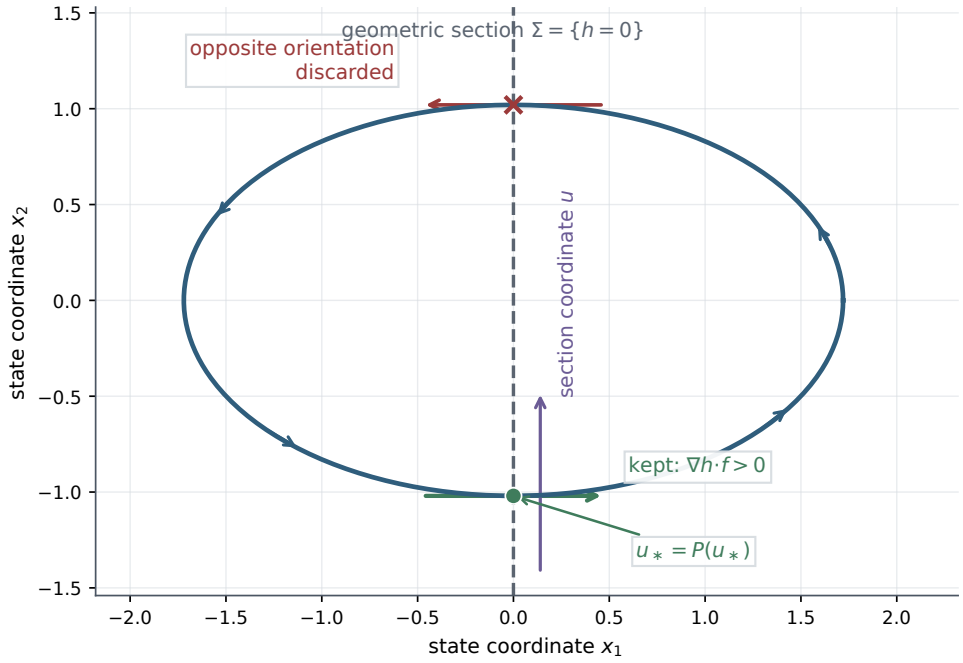


Figure 2.4: An oriented Poincaré section compares the internal state at repeated crossings. For the illustrated simple oscillation, retaining the same direction excludes the half-cycle crossing. A periodic trajectory gives a fixed point of the map between these events.

Starting there and keeping that direction requires a full circuit before the next observation. The section coordinate then agrees with its initial value: the map has a fixed point.

For 2SEG, the observation will be $\delta = 0$ with $\chi\sigma > 0$, where $\chi = \pm 1$ selects direction. At two such crossings, the effective potential and inertia are the same. If the effective oscillator also has the same energy, its joint speed has the same magnitude; the crossing direction fixes its sign. This is the physical reason a scalar energy condition can determine the remaining joint-velocity condition. The next chapter states exactly which other variables must already agree and proves the resulting equivalence.

For a more complicated motion, the number of section crossings and the least positive period must be examined: a recorded integration interval can contain several primitive cycles. This matters for interpreting the three-segment illustrations later.

2.5.2 The Natural Locomotion Manifold and a speed coordinate

One natural cycle is one closed internal-state trajectory. A regular one-parameter family of neighbouring natural cycles, swept through all their phases, forms locally a two-dimensional invariant surface: the *Natural Locomotion Manifold* (NLM). A representation is

$$z = Z(a, \varphi), \qquad \dot{a} = 0, \qquad \dot{\varphi} = \frac{2\pi}{T(a)}. \tag{2.33}$$

The coordinate a selects the cycle and φ locates the state within that cycle. This is the same distinction made visually in Figure 2.2: choosing an orbit and advancing along it are different operations. For the constrained locomotor, however, the family label a need not be canonical action.

A particularly informative cycle label is mean forward body speed,

$$\bar{v}(a) = \frac{1}{T(a)} \int_0^{T(a)} v(t; a)\, \mathrm{d}t. \tag{2.34}$$

Where $\mathrm{d}\bar{v}/\,\mathrm{d}a \neq 0$, the inverse function theorem makes $(\bar{v}, \varphi)$ local coordinates on the same regular surface. One can then think of the NLM as a repertoire organized by average advance: different cycles give different mean forward speeds. The current two-segment computations test this condition along their sampled connected branch.

This picture does not say that the passive robot changes cycles by itself. Under the fixed conservative dynamics, $\dot{a} = 0$; changing the selected cycle requires an intervention. It also does not identify $\bar{v}$ with world-axis displacement divided by period. If the robot turns, $\Delta x/T$ depends on heading as well as forward body speed.

The regularity qualification has a mathematical role. After fixing the phase, independent periodic-orbit equations of full derivative rank give a local solution branch; restoring phase gives a surface when the phase direction and the family direction are independent. A fold of a plotted speed or energy need not destroy the branch, but it can make that observable unsuitable as a coordinate. Appendix G gives the precise local statements. Numerical continuation, developed after the energy principle in Chapter 4, computes neighbouring solutions and tests this geometry.

2.6 From an oscillator's energy cycle to a locomotor's energy cycle

The ingredients can now be put together. An ordinary conservative oscillator repeats as energy moves between kinetic and potential forms. In a locomotor, the effective oscillatory and propulsive energies provide another partition, entirely within the internal mechanics. The effective oscillator itself still has kinetic and potential energy.

For the one-joint locomotor, the effective oscillator has one position and one momentum, just as the simple pendulum does. Its regular bounded energy levels therefore have the action–angle organization already described. With two joints, the effective oscillator instead has four state coordinates, as in the double pendulum, and one energy value does not select one coherent motion. Table 2.2 makes this scalar/vector contrast explicit before the locomotor energy partition is derived.

Table 2.2: The oscillator analogy identifies what carries over and what locomotion adds. It compares effective oscillatory dimensions, not the total degrees of freedom of the mechanisms.

	One coordinate	Several coordinates
Natural oscillation	A regular bounded energy curve is a closed phase-plane orbit.	A modal family selects coherent cycles within a higher-dimensional energy surface.
Natural locomotion	The 2SEG theorem relates zero cycle-integrated POE to the remaining joint-speed condition on a section.	For 3SEG, one energy condition is insufficient; all internal variables must repeat.
Added physical requirement	The internal motion solves the locomotor dynamics and accumulates a non-trivial pose increment, with translation for the propulsive cases.	

The next chapter derives these two energy channels and the *propulsion–oscillator exchange* (POE). Its principle states that natural locomotion requires the energy transferred between the effective oscillatory and propulsive channels to be returned over a cycle, without average gain or loss in either channel. The scalar 2SEG case makes this intuitive rebound a precise characterization under stated section hypotheses. Computation and mechanical design then use that structure.

After the conservative families have been identified, dissipation will change the energy budget. Without energy input, persistent losses cannot sustain the same locomotion. An actuator must supply work; its timing can accompany the natural organization, as pushing a swing at the appropriate phase does. The distinction between internal POE and external input–loss balance is retained when Chapter 6 introduces resonant maintenance and Chapter 7 develops feedback.

Part II

An energy principle and natural gait families

Chapter 3

An energy-exchange principle for natural locomotion

A pendulum gives back the energy it stores: kinetic energy becomes potential energy, then potential energy becomes kinetic energy again. A locomotor asks for an additional exchange. Its joints oscillate while the body moves through the environment; the mechanical coupling between these motions redistributes energy even when their total remains constant. Can that redistribution repeat without progressively emptying one part of the motion into another?

Our principle is that, for locomotion to be natural, the energy exchanged between an *effective oscillatory channel* and a *propulsive channel* must balance over a cycle, with no mean gain or loss in either channel. We call this transfer *propulsion–oscillator exchange* (POE). The analogy puts the two pairs of energies face to face, but does not identify them: the effective oscillator itself contains kinetic and potential energy. Both locomotor channels are defined within the same internal dynamics, not on opposite sides of a division between internal and external motion.

This chapter gives the principle its mathematical content. First, kinetic coupling and potential energy define the two reservoirs. Their rates expose the power redistributed through the environmental constraints. For the two-segment locomotor, balanced POE energy then characterizes natural cycles exactly within a regular return setting. With the body-velocity conditions, this is a necessary-and-sufficient description of the cycles generating the local Natural Locomotion Manifold, rather than merely a necessary energy check. The next chapter turns that characterization into a calculation.

3.1 Two motions within one internal state

Consider again the two-segment robot of the opening example. Its relative joint angle δ changes its shape; the pose $g = (x, y, \theta) \in \mathrm{SE}(2)$ places it in the plane. In this application the environment is represented by an ideal lateral velocity constraint. The body can translate along its heading and rotate, with velocities v and ω. Together with the joint rate $\sigma = \dot{\delta}$, these give the internal state

$$z = (\delta, v, \omega, \sigma). \tag{3.1}$$

Neither v nor ω is a pose coordinate: both participate in the internal dynamics after pose has been removed by symmetry. The pose is recovered through

$$\dot{x} = v\cos\theta, \qquad \dot{y} = v\sin\theta, \qquad \dot{\theta} = \omega. \tag{3.2}$$

A nonconstant solution may repeat all four components of z after a time T yet accumulate a translation. We seek that repeated internal motion with an open pose, not a closed path of the robot in the plane. Constant straight translation is a useful reference, but has no internal oscillatory cycle.

The energetic argument can be developed before choosing numerical masses or stiffnesses. We work with an autonomous, unforced, conservative mechanical model whose dynamics do not depend on absolute planar pose. Its potential is a function of shape. Elasticity is one way to obtain such a potential; in other settings gravity or another conservative interaction may provide it. The ideal environmental constraints supply no total work. The derivation below concerns smooth, continuously active velocity constraints and nonsingular admissible coordinates; the common constrained-mechanics appendix specifies this setting.

Thus the physical motivation is general, while the identities and theorems are established for a declared conservative model. Losses and actuation enter separately in Chapter 6.

3.2 Extracting the effective oscillator from the mechanics

3.2.1 What completing a square means physically

Write the admissible velocity as $\nu = (c, w)$, where c contains body velocities and $w = \dot{\boldsymbol{\delta}}$ contains shape rates. For 2SEG, $\boldsymbol{\delta} = \delta$, $w = \sigma$, and $c = (v, \omega)^{\mathsf{T}}$. The kinetic energy is

$$K = \frac{1}{2}\begin{bmatrix} c \\ w \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} M_{cc}(\boldsymbol{\delta}) & M_{cw}(\boldsymbol{\delta}) \\ M_{cw}(\boldsymbol{\delta})^{\mathsf{T}} & M_{ww}(\boldsymbol{\delta}) \end{bmatrix} \begin{bmatrix} c \\ w \end{bmatrix}, \qquad M(\boldsymbol{\delta}) > 0. \tag{3.3}$$

The off-diagonal block expresses kinetic coupling: joint motion changes how body motion contributes to energy, and conversely. Assigning every term involving a body velocity to one reservoir would leave the mixed terms unresolved.

Instead, hold shape and shape rate fixed and ask which body velocity minimizes kinetic energy. As in a quadratic polynomial, its derivative must vanish:

$$M_{cc}c_\star + M_{cw}w = 0, \qquad c_\star = -M_{cc}^{-1}M_{cw}w, \qquad q_c = c - c_\star. \tag{3.4}$$

The coordinate q_c measures the difference from that minimizing motion. It is not the body velocity itself: $q_c = 0$ can have $c = c_\star \neq 0$. Completing the square gives

$$M_\perp = M_{ww} - M_{cw}^{\mathsf{T}}M_{cc}^{-1}M_{cw}, \tag{3.5}$$

$$E_\perp(\boldsymbol{\delta}, w) = \frac{1}{2}w^{\mathsf{T}}M_\perp(\boldsymbol{\delta})w + V(\boldsymbol{\delta}), \tag{3.6}$$

$$E_{\mathrm{prop}}(\boldsymbol{\delta}, c, w) = \frac{1}{2}q_c^{\mathsf{T}}M_{cc}(\boldsymbol{\delta})q_c, \qquad E = E_\perp + E_{\mathrm{prop}}. \tag{3.7}$$

The matrix $M_\perp$ is a Schur complement, positive definite because M is. It is the effective inertia felt in the oscillatory channel after the minimizing body motion has been accounted for.

Figure 3.1 puts the two energy partitions side by side. In panel (a), a pendulum repeatedly converts kinetic energy into potential energy and back. Panel (b) applies the cyclic-exchange idea to the locomotor, but with different reservoirs: the effective oscillator already stores both potential and effective kinetic energy. The propulsive reservoir contains the complementary body-motion contribution. It is not all the kinetic energy of the translating or rotating body: part of that motion is already included in $c_\star$ and $M_\perp$. Both reservoirs depend on the internal state. Their split follows the declared body–shape partition; it is not a unique coordinate-free attribution of which body caused a given displacement.

3.2.2 Specialization to the two-segment robot

For the velocity order (v, ω, σ) the normalized kinetic matrix is

$$M(\delta) = \begin{bmatrix} 1 & -\rho & -\rho \\ -\rho & B_{11} & B_{12} \\ -\rho & B_{12} & B_{22} \end{bmatrix}. \tag{3.8}$$

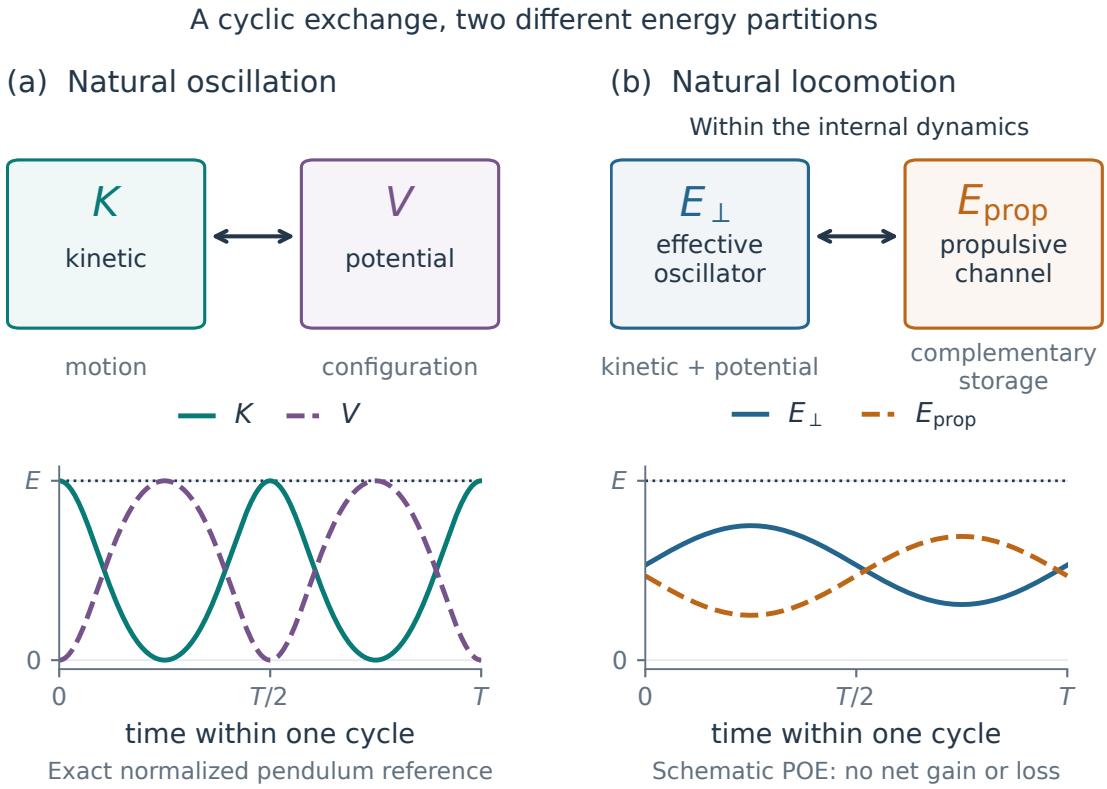


Figure 3.1: An energy analogy, not an identification of energy partitions. (a) Kinetic and potential energy along an exact normalized pendulum libration of amplitude 1.8 rad. (b) Schematic redistribution between the effective oscillatory and propulsive channels of a locomotor. Each plot is scaled by its constant total energy E; both channels regain their initial allocation after a cycle. The right-hand curves illustrate the POE principle, not a computed robot trajectory. Ideal environmental constraints redirect energy without providing an external supply. The effective oscillator itself contains kinetic and potential energy.

The B coefficients represent angular inertia and coupling; ρ is the translation–rotation coupling induced by joint deflection. They follow from the masses, lengths and inertias in Appendix B; the next chapter gives the numerical design. The body block and mixed column are

$$M_{cc} = \begin{bmatrix} 1 & -\rho \\ -\rho & B_{11} \end{bmatrix}, \qquad m_{c\sigma} = \begin{bmatrix} -\rho \\ B_{12} \end{bmatrix}, \qquad \Delta = \det M_{cc} = B_{11} - \rho^2. \tag{3.9}$$

The scalar effective inertia is therefore

$$M_\perp(\delta) = B_{22} - \frac{B_{12}^2 + \rho^2(B_{11} - 2B_{12})}{\Delta}, \tag{3.10}$$

and its oscillatory energy is

$$\boxed{E_\perp(\delta, \sigma) = \frac{1}{2} M_\perp(\delta)\sigma^2 + U(\delta).} \tag{3.11}$$

The spring potential U belongs to this particular robot. Neither $M_\perp$ nor $E_\perp$ is obtained by discarding body inertia: both incorporate its coupling. The effective oscillator is consequently a mechanical object, not the visibly oscillating link considered alone.

3.3 An auxiliary natural oscillation and its action

3.3.1 Closing the coupling changes the dynamical model

The energy in Equation (3.6) defines the auxiliary Lagrangian

$$L_\perp(\boldsymbol{\delta}, w) = \frac{1}{2} w^\mathsf{T} M_\perp(\boldsymbol{\delta}) w - V(\boldsymbol{\delta}). \tag{3.12}$$

Its Euler–Lagrange equations describe a conservative oscillator. We call this operation *closing the propulsive coupling*. It isolates a problem in which effective-oscillator energy is conserved and natural-oscillation tools apply. It does not physically immobilize the body, and does not assume that $q_c = 0$ is an invariant subset of the original locomotor.

For one coordinate the auxiliary equation is

$$\dot{\delta} = \sigma, \qquad \dot{\sigma} = F_{\mathrm{cl}}(\delta, \sigma) = -\frac{U'(\delta) + \frac{1}{2} M'_\perp(\delta)\sigma^2}{M_\perp(\delta)}. \tag{3.13}$$

The effective-inertia derivative is essential: it follows from the Euler–Lagrange equation. Substitution gives $\dot{E}_\perp = 0$. Regular oscillations between turning points δ_- and δ_+ have

$$\sigma_\pm(\delta; \mathcal{E}) = \pm\sqrt{\frac{2\{\mathcal{E} - U(\delta)\}}{M_\perp(\delta)}}, \qquad U(\delta_\pm) = \mathcal{E}. \tag{3.14}$$

These are natural oscillations of the auxiliary model. They provide starting points for the complete locomotor, whose coupling can change timing, waveform and energy allocation. A closed auxiliary curve is not already a locomotion solution. Thus the oscillator gives us an organized starting motion; the next question is how much energy the original coupling moves into and out of it.

3.3.2 Action connects modal geometry and projected momentum

One coordinate selects a regular auxiliary energy curve, and another locates a state on it. Its canonical momentum and action are

$$p_\perp = \frac{\partial L_\perp}{\partial \sigma} = M_\perp(\delta)\sigma, \qquad J_{\mathrm{cl}}(\mathcal{E}) = \frac{1}{2\pi} \oint p_\perp \, \mathrm{d}\delta. \tag{3.15}$$

The action is the enclosed momentum–angle area divided by 2π. The angle variable advances once around the oscillation; it is a phase, not generally the physical joint angle δ.

This quantity also connects directly to the momenta of constrained mechanics. Define

$$\begin{bmatrix} \pi_c \\ \pi_w \end{bmatrix} = \frac{\partial K}{\partial (c, w)} = M(\boldsymbol{\delta}) \begin{bmatrix} c \\ w \end{bmatrix}. \tag{3.16}$$

These projected momenta are often called pseudo-momenta because the admissible velocity basis need not be a coordinate basis. The block relations imply

$$\begin{aligned} \pi_c &= M_{cc} c + M_{cw} w = M_{cc} q_c, \\ \pi_w - M_{cw}^\mathsf{T} M_{cc}^{-1} \pi_c &= M_\perp w. \end{aligned} \tag{3.17}$$

The 2SEG action integrand is thus the effective transverse component of projected momentum after the body contribution is eliminated, not the entire pseudo-momentum vector. Its algebra links the oscillator's action to the same mechanics used for locomotion reduction (Bloch, 2015; Arnold, 1989).

Period and action follow from one-dimensional quadratures:

$$T_{\mathrm{cl}}(\mathcal{E}) = 2\int_{\delta_-}^{\delta_+} \sqrt{\frac{M_\perp(\delta)}{2\{\mathcal{E} - U(\delta)\}}}\,\mathrm{d}\delta, \tag{3.18}$$

$$J_{\mathrm{cl}}(\mathcal{E}) = \frac{1}{\pi}\int_{\delta_-}^{\delta_+} \sqrt{2M_\perp(\delta)\{\mathcal{E} - U(\delta)\}}\,\mathrm{d}\delta. \tag{3.19}$$

On a regular oscillatory annulus, $\mathrm{d}J_{\mathrm{cl}}/\,\mathrm{d}\mathcal{E} = T_{\mathrm{cl}}/(2\pi)$, so its angular frequency is $\mathrm{d}\mathcal{E}/\,\mathrm{d}J_{\mathrm{cl}}$. A separatrix or degenerate turning point requires a different description.

The pendulum analogy is now precise. The complete 2SEG robot has a pose, a joint and three admissible velocities, yet the auxiliary construction exposes a one-coordinate oscillator with closed energy curves and an action. This does not Hamiltonize the full nonholonomic locomotor. Its effective energy will generally vary when the original coupling is restored.

3.4 Power redirected between the two channels

Let $F_{\mathrm{op}}(\boldsymbol{\delta}, c, w)$ be the shape acceleration of the original unforced locomotor, and $F_{\mathrm{cl}}(\boldsymbol{\delta}, w)$ the auxiliary Euler–Lagrange acceleration. At the same internal state define

$$Q_{\mathrm{POE}} = M_\perp(F_{\mathrm{op}} - F_{\mathrm{cl}}), \qquad P_{\mathrm{POE}} = w^{\mathsf{T}} Q_{\mathrm{POE}}. \tag{3.20}$$

The generalized force measures the difference between coupled and auxiliary shape dynamics. It is not an actuator added to the passive model. Differentiating $E_\perp$ along the original flow, the terms due to potential and shape-dependent inertia cancel those of the auxiliary Euler–Lagrange equation. Therefore

$$\boxed{\dot{E_\perp} = P_{\mathrm{POE}}, \qquad \dot{E_{\mathrm{prop}}} = -P_{\mathrm{POE}}, \qquad \dot{E} = 0.} \tag{3.21}$$

Positive POE power means that the effective oscillator receives energy; negative power means it gives energy to the propulsive channel. The ideal constraints add no reservoir: together with inertia, they redirect energy between the two mechanical motions.

For 2SEG the identity can be checked explicitly. Solving the original mechanical equations gives the joint acceleration

$$F_{\mathrm{op}} = \frac{r_q(\delta)\{v\omega - \rho(\delta)(\sigma+\omega)^2\} - U'(\delta)}{M_\perp(\delta)}, \tag{3.22}$$

where $r_q = A_q/\Delta$, with A_q given in Appendix C. Comparing it with Equation (3.13) gives

$$M_\perp(F_{\mathrm{op}} - F_{\mathrm{cl}}) = r_q\{v\omega - \rho(\sigma+\omega)^2\} + \frac{1}{2}M_\perp'\sigma^2. \tag{3.23}$$

Multiplication by σ is exactly the derivative of Equation (3.11) along the opened field. Setting only ω or $v\omega$ to zero in the original acceleration would not produce the auxiliary oscillator: its effective-inertia derivative must remain.

Integration over any regular interval gives

$$\boxed{I_{\mathrm{POE}}(T) = \int_0^T P_{\mathrm{POE}}(t)\,\mathrm{d}t = E_\perp(T) - E_\perp(0).} \tag{3.24}$$

The integral is an energy, whereas P_{POE} is a power. It is defined before a periodic solution has been found, and can be tested on a candidate that does not repeat. An internal cycle necessarily has $I_{\mathrm{POE}}(T) = 0$.

This is the balanced-rebound statement: each channel regains the energy it has ceded over the complete cycle. Zero net transfer does not require zero instantaneous power; the energy can pass in both directions. Conversely, the identity alone does not force every motion to have a strictly nonzero transfer. Nonconstant internal motion, full dynamical repetition and displacement remain part of natural locomotion, not consequences of the zero of an arbitrary integral.

3.5 Why one effective coordinate gives an exact criterion

3.5.1 Observe the joint at a reproducible instant

Imagine marking each instant at which the joint crosses $\delta = 0$ in the same direction. This Poincaré section compares states at the same stage of the oscillation. The event fixes the angle, and its orientation distinguishes the two directions in which the oscillating joint crosses zero.

Three quantities remain: the joint speed and the two body velocities. Suppose the latter recover their initial values. Does the joint recover its velocity too? On this section the effective inertia and potential have the same values at both events. Recovering effective-oscillator energy then fixes the magnitude of the joint velocity, and the crossing direction fixes its sign. This is why the energy principle becomes sufficient in the scalar setting.

Choose $\chi \in \{-1, +1\}$ and set

$$\Sigma^{\chi} = \{z : \delta = 0,\ \chi\sigma > 0\}, \qquad \zeta = \chi\sigma > 0. \tag{3.25}$$

Here χ selects the crossing direction, independently of the direction of mean translation. At consecutive regular crossings z_0, z^+, the section energy is

$$E_{\Sigma}(\zeta) = \frac{1}{2} M_{\perp}(0)\zeta^2 + U(0), \qquad \frac{\mathrm{d}E_{\Sigma}}{\mathrm{d}\zeta} = M_{\perp}(0)\zeta > 0. \tag{3.26}$$

One effective-energy value selects one positive crossing speed. Read Figure 3.2 from left to right. The two marked events have the same joint angle and crossing direction, but a candidate motion need not yet have the same speed. Their different heights on the section become different energies on the strictly increasing curve. Making the POE energy zero removes exactly this speed difference. The body-velocity conditions then complete internal periodicity.

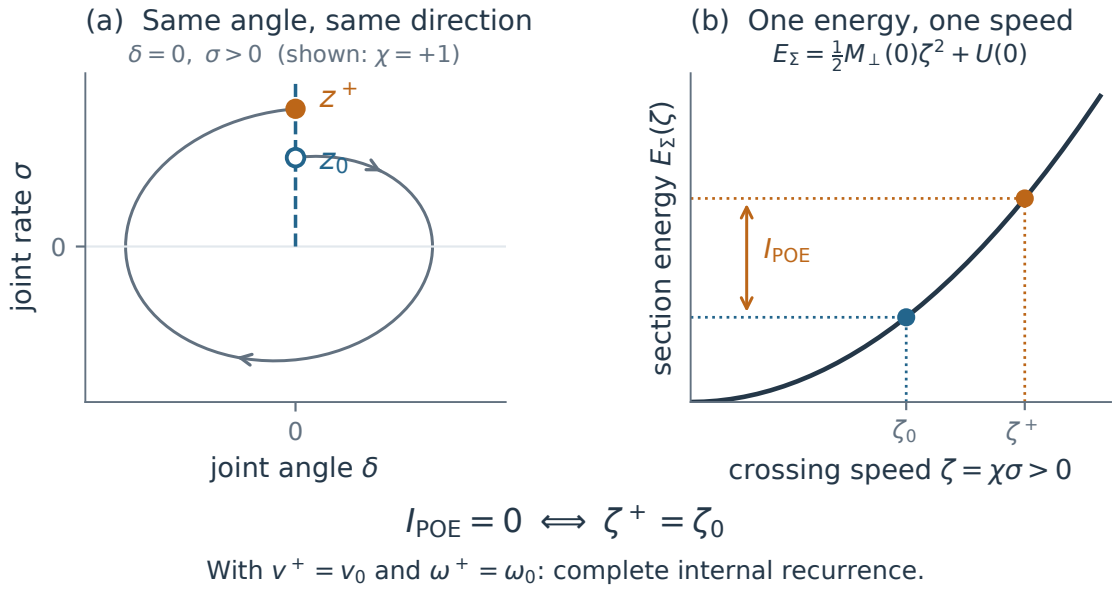


Figure 3.2: Why one energy condition determines the remaining joint speed. (a) A conceptual angle–rate projection of a candidate motion, with two same-direction section crossings. The illustrated candidate does not close. (b) On the positive crossing coordinate ζ, the section energy is strictly increasing; the vertical difference is I_{POE}. Equality of these energies therefore fixes the joint rate. Repetition of the two body velocities is an additional condition for full internal recurrence. The diagram is not a numerical trajectory or an assertion that the coupled motion follows an auxiliary energy loop.

3.5.2 Statement and proof

Assumption 3.1 (Regular scalar crossing)**.** The trajectory follows the original unforced conservative 2SEG field and remains in nonsingular admissible coordinates with $M_\perp > 0$. It has a finite next transverse crossing of the same oriented section Σ^χ. The storage-rate identity Equation (3.21) holds along it.

Theorem 3.2 (Scalar POE–return equivalence)**.** *Under Assumption 3.1, the POE energy between the crossings is*

$$I_{\mathrm{POE}} = \frac{M_\perp(0)}{2}(\zeta^+ + \zeta_0)(\zeta^+ - \zeta_0). \tag{3.27}$$

Consequently,

$$\boxed{I_{\mathrm{POE}} = 0 \quad \Longleftrightarrow \quad \zeta^+ = \zeta_0.} \tag{3.28}$$

Proof. By Equation (3.24), I_{POE} is the difference of endpoint effective-oscillator energies. Both endpoints have $\delta = 0$, so their potentials cancel and their effective inertias coincide. With $\sigma = \chi\zeta$,

$$I_{\mathrm{POE}} = \frac{1}{2}M_\perp(0)\{(\zeta^+)^2 - \zeta_0^2\}.$$

Factor the difference of squares. The prefactor $M_\perp(0)(\zeta^+ + \zeta_0)/2$ is strictly positive. Hence the energy difference vanishes exactly when the oriented joint speed repeats. □

Corollary 3.3 (Characterization of the natural 2SEG cycles)**.** *In the same regular setting, impose $v^+ = v_0$ and $\omega^+ = \omega_0$. Then $I_{\mathrm{POE}} = 0$ is necessary and sufficient for repetition of the complete internal state. The resulting nonconstant motion is a propulsive natural cycle when its separately reconstructed pose has nonzero translation. For a regular one-parameter family of these cycles with an embedded phase sweep, these conditions characterize the corresponding local NLM.*

Proof. The section event supplies the repeated joint angle; the two imposed equalities supply the body velocities; the theorem supplies the remaining joint rate. All components of Equation (3.1) therefore repeat. Conversely, an internal cycle repeats its effective energy and has $I_{\mathrm{POE}} = 0$. Reconstruction distinguishes a locomoting cycle from a motion with closed pose. Restoring phase on a regular embedded one-parameter family gives its local two-dimensional invariant surface, as developed in Appendix G. □

The family statement is conditional on regularity, not a claim that every robot design admits natural locomotion. Its positive content is precise: the energy principle and internal periodicity are equivalent on the stated 2SEG section once the body-velocity conditions are included. It is not merely a check applied after every coordinate has already been made to repeat.

3.5.3 What this makes possible

Instead of comparing the two joint rates directly, a calculation may seek

$$\mathcal{R}_{\mathrm{phys}}(z_0) = \begin{bmatrix} I_{\mathrm{POE}} \\ v^+ - v_0 \\ \omega^+ - \omega_0 \end{bmatrix} = 0 \quad \text{on } \Sigma^{\chi}. \tag{3.29}$$

An additional label selects the wanted cycle. Total energy conservation can make return conditions locally dependent, so the POE equation must not be counted as extra information. Chapter 4 develops the solve and its numerical realization. Its finite-speed coordinates facilitate calculation; they are not needed for the physical crossing argument.

Orientation is essential. Without it, σ and $-\sigma$ have the same effective energy. A zero POE balance between opposite crossings need not recover the joint velocity. Likewise, the POE equation alone leaves the body velocities unspecified. These examples explain the hypotheses while preserving the full strength of the result under them.

3.6 Where the scalar argument stops

The energy split and auxiliary Lagrangian extend to more shape coordinates; the decisive injectivity does not. With two joints, the same effective energy can be distributed among several joint rates and configurations. On a regular higher-dimensional section,

$$E_{\Sigma}(y^+) = E_{\Sigma}(y_0) \tag{3.30}$$

places y^+ on a level set rather than selecting its endpoint uniquely. Further independent conditions must make all internal components repeat. Chapter 5 demonstrates this distinction with two recognizable joint organizations of the same robot.

The principle therefore expresses cyclic energy allocation in the conservative model and, for 2SEG, gives an exact characterization within the regular section setting. For the two-joint robot it remains necessary, while the calculation resolves vector periodicity as well. Conservation of total energy cannot replace that information: total energy is constant even on unforced motions that do not repeat. We now use the scalar characterization to calculate a repertoire of motions and examine its dependence on the robot's mechanics.

Chapter 4

Computing a repertoire of natural gaits

A criterion for natural locomotion becomes useful to a roboticist when it leads to an actual motion: joint angles and rates, body velocities, a period and a displacement. The POE principle does not prescribe those signals. It identifies a mechanical condition that they must satisfy. This chapter shows how to calculate them, first for one robot and then while changing a mechanical parameter.

The two-segment locomotor is a useful anchor because its one effective coordinate makes the analytical result of Chapter 3 exact. Its complete internal dynamics still has four state coordinates. The method starts from an organized oscillation, restores the original coupling and adjusts an initial state and a period until the complete internal motion repeats. Continuing to neighbouring solutions then produces a repertoire, rather than one trajectory to be followed by a controller.

We use the mechanism introduced in Rajaomarosata et al. (2025a). The numerical results distinguish 29 independently recovered scalar roots from a connected 49-node continuation at unchanged mechanical parameters. Mean forward speed provides a physical local coordinate for that latter family: it selects the gait, while phase locates the current state on it. The second part of the chapter changes inertia instead. It derives an exact mechanical relation and compares its small-oscillation prediction with finite-amplitude natural cycles of different robot designs.

4.1 Mechanical model

4.1.1 Configuration, constraint, and pose reconstruction

The 2SEG locomotor consists of a planar chassis and one internal rigid segment joined by a torsional spring. Its configuration is

$$q = (g, \delta), \qquad g = (x, y, \theta) \in \mathrm{SE}(2), \tag{4.1}$$

where δ is the relative joint angle. The chassis can roll along its heading but cannot move sideways at its contact point. Figure 4.1 fixes the positive directions and reduced variables before the constraint is written.

$$A(\theta)\dot{q} = 0, \qquad A(\theta) = \begin{bmatrix} -\sin\theta & \cos\theta & 0 & 0 \end{bmatrix}. \tag{4.2}$$

This is called a *nonholonomic* constraint: it restricts instantaneous velocities and cannot, in general, be integrated into a fixed relation among x, y, and θ. An admissible velocity can be written as

$$\dot{q} = H(\theta)\eta, \qquad H(\theta) = \begin{bmatrix} \cos\theta & 0 & 0 \\ \sin\theta & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \qquad \eta = \begin{bmatrix} v & \omega & \sigma \end{bmatrix}^{\mathsf{T}}, \tag{4.3}$$

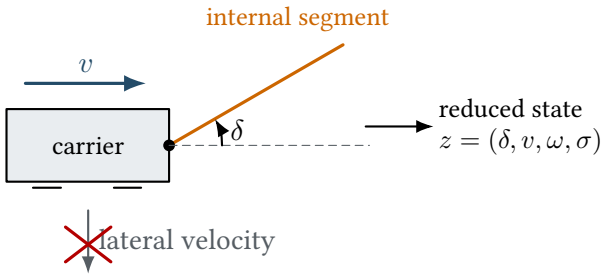


Figure 4.1: The two-segment locomotor and its reduced variables. The ideal wheel-like constraint removes lateral contact velocity; it does not remove forward speed or yaw. The spring-mounted internal segment is described by joint angle δ and joint rate σ.

with longitudinal speed v, yaw rate $\omega = \dot{\theta}$, and joint rate $\sigma = \dot{\delta}$. Because $A(\theta)H(\theta) = 0$, this representation enforces the lateral constraint by construction (Bloch, 2015; Rajaomarosata et al., 2025a).

Task-oriented models of nonholonomic mobile manipulators likewise parameterize available velocities to expose task and redundant directions (De Luca et al., 2007). Here the same kernel parameterization serves a different, dynamic purpose: it projects the Lagrange–d'Alembert equations and eliminates the ideal reaction forces.

This is the first application of the constrained-reduction pattern used again for the three-segment locomotor. Starting from the Lagrange–d'Alembert equation with reaction $A(q)^{\mathsf{T}}\lambda$, multiplication by $H(q)^{\mathsf{T}}$ gives

$$H(q)^{\mathsf{T}}\left[\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q\right] = 0, \qquad \operatorname{Im} H(q) = \ker A(q). \tag{4.4}$$

The reaction disappears because every column of H is an admissible velocity. Equivalently, an ideal reaction has zero power on every admissible motion. It does not energize the locomotor; the constraint distribution and the kinetic metric instead determine how its admissible velocity components are coupled.

Table 4.1 makes the dimensions explicit. The same four operations—constraint, kernel basis, projected dynamics, and later pose reconstruction—will be used with the matrix N in Chapter 5.

The reusable scope is precise: the projection applies to smooth, bilateral, ideal, continuously active, constant-rank velocity constraints. It does not by itself describe impacts, unilateral or switching

Table 4.1: Common projected Lagrange–d'Alembert pattern specialized to the 2SEG locomotor. The effective transverse dimension is obtained only after the two carrier velocities have subsequently been eliminated; it is not the dimension of the complete locomotor.

Operation	General object	2SEG realization
configuration	$q \in Q$	$q = (x, y, \theta, \delta)$: three pose coordinates and one joint
velocity constraint	$A(q)\dot{q} = 0$	one lateral no-slip row, of rank one
admissible basis	$\dot{q} = H(q)\eta$, $\operatorname{Im} H = \ker A$	$\eta = (v, \omega, \sigma)$, three admissible velocities
projected mechanics	project the Lagrange–d'Alembert balance with H^{T}	three reduced acceleration balances
transverse oscillator	carrier velocities eliminated from the kinetic energy	one effective transverse degree of freedom (δ, σ)
pose output	integrate the group equation after the reduced solve	reconstruct (x, y, θ) from (v, ω)

contacts, or stick–slip dynamics.

The distinction between exact reduced return and net pose displacement was introduced in Figure 1.1.

The dynamics are solved in the four-dimensional state

$$z = (\delta, v, \omega, \sigma)^{\mathsf{T}}. \tag{4.5}$$

The planar pose is then reconstructed from

$$\dot{x} = v\cos\theta, \qquad \dot{y} = v\sin\theta, \qquad \dot{\theta} = \omega. \tag{4.6}$$

Consequently, a closed trajectory in z need not close in (x, y, θ). That distinction is the elementary mechanism by which exact reduced return and locomotion coexist.

4.1.2 Compact reduced dynamics

The ratios α and β describe the two geometric lever arms relative to their sum, μ is the internal segment's mass fraction, and j_1, j_2 are inertias normalized by total mass times squared reference length. The dimensional definitions and reference clock are given in appendix B. The computations select a design class in which one fraction ε sets the length and mass split, while γ scales the inertias:

$$\alpha = \varepsilon, \qquad \beta = \mu = 1 - \varepsilon, \qquad j_1 = \gamma\alpha^3, \qquad j_2 = \gamma\beta^3, \tag{4.7}$$

and the passive joint potential

$$U(\delta) = \frac{1}{2}k_2\delta^2 + \frac{1}{4}k_4\delta^4. \tag{4.8}$$

The reference values are

$$\varepsilon = 0.681, \qquad \gamma = 2.504, \qquad k_2 = 1, \qquad k_4 = 10. \tag{4.9}$$

Applying Equation (4.4) to the kinetic and potential energies of this locomotor gives the compact balance

$$M(\delta)\dot{\eta} + c(\delta, \eta) + \begin{bmatrix} 0 & 0 & U'(\delta) \end{bmatrix}^{\mathsf{T}} = 0. \tag{4.10}$$

Here c contains the velocity-quadratic inertial terms and

$$M(\delta) = \begin{bmatrix} 1 & -\rho & -\rho \\ -\rho & B_{11} & B_{12} \\ -\rho & B_{12} & B_{22} \end{bmatrix}, \tag{4.11}$$

with

$$\rho(\delta) = \mu\beta\sin\delta, \tag{4.12a}$$

$$B_{11}(\delta) = j_1 + j_2 + \alpha^2 + \mu\beta^2 + 2\mu\alpha\beta\cos\delta, \tag{4.12b}$$

$$B_{12}(\delta) = j_2 + \mu\beta^2 + \mu\alpha\beta\cos\delta, \tag{4.12c}$$

$$B_{22} = j_2 + \mu\beta^2. \tag{4.12d}$$

The dimensional derivation, the explicit inertial vector, the zero-work check, and the scaling dictionary are given in appendix B; that appendix uses the same velocity basis and points back to the present reduced equations. In physical terms, ρ couples translation to the two angular rates when the joint is deflected, while the terms proportional to $v\omega$ couple a moving chassis to its yaw and internal dynamics.

Where $M(\delta)$ is nonsingular, define

$$a(\delta,\eta) = -M(\delta)^{-1}\left(c(\delta,\eta) + \begin{bmatrix}0 & 0 & U'(\delta)\end{bmatrix}^{\mathsf{T}}\right) = \begin{bmatrix}a_v & a_\omega & a_\sigma\end{bmatrix}^{\mathsf{T}}. \tag{4.13}$$

The reduced dynamics are then

$$\dot{z} = f(z) = \begin{bmatrix}\sigma & a_v & a_\omega & a_\sigma\end{bmatrix}^{\mathsf{T}}. \tag{4.14}$$

With no forcing or damping, the mechanical energy

$$E(z) = \frac{1}{2}\eta^{\mathsf{T}} M(\delta)\eta + U(\delta) \tag{4.15}$$

is conserved.

4.2 From an organized oscillation to a passive cycle

4.2.1 A moving equilibrium is not an oscillatory gait

For every constant longitudinal velocity s_0, the state

$$z_{\mathrm{re}}(s_0) = (0, s_0, 0, 0)^{\mathsf{T}}, \qquad f(z_{\mathrm{re}}(s_0)) = 0 \tag{4.16}$$

is an internal equilibrium. Its reconstructed motion is

$$x = x_0 + s_0 t\cos\theta_0, \quad y = y_0 + s_0 t\sin\theta_0, \quad \theta = \theta_0. \tag{4.17}$$

It is therefore a relative equilibrium: the robot translates even though its internal state does not change. Assigning it an arbitrary period would not create a natural oscillatory cycle.

This moving reference also explains why periodicity is a problem with structure. There is a family of equilibria, not one isolated equilibrium, and its speed tangent is neutral:

$$Df(z_{\mathrm{re}}(s_0))(0,1,0,0)^{\mathsf{T}} = 0. \tag{4.18}$$

Oscillations are sought transverse to this tangent. With $\xi_\perp = (\delta_1, \omega_1, \sigma_1)^{\mathsf{T}}$, linearization gives

$$\dot{\xi}_\perp = J_\perp(s_0)\xi_\perp. \tag{4.19}$$

For the reference design at zero speed, the oscillatory eigenpair is $\lambda_\pm = \pm 3.30569i$. At nonzero speed the corresponding pair need not remain exactly imaginary. Its eigenvector still supplies an organized initial guess, but not a theorem asserting a neighbouring periodic orbit.

For an eigenvalue $\lambda = a_\lambda + i\Omega$ and a complex eigenvector $p = p_r + ip_i$, embed it as $\widehat{p} = (p_\delta, 0, p_\omega, p_\sigma)^{\mathsf{T}}$. A small seed is

$$z_{\mathrm{seed}}(t) = z_{\mathrm{re}}(s_0) + A\{\widehat{p}_r\cos\Omega t - \widehat{p}_i\sin\Omega t\}. \tag{4.20}$$

Nonlinear correction determines whether this guess can become a passive cycle. The final variation of v about its mean, its waveform and its period are not fixed by this harmonic seed.

4.2.2 Use the effective oscillator without mistaking it for the robot

The auxiliary energy curves of Equation (3.14) provide a second, finite-amplitude route to an organized initial guess. Their action and period are available by quadrature. A point (δ, σ) on a chosen auxiliary oscillation can be augmented with body velocities, for example by the minimizing velocity $c_\star$ of Equation (3.4). The result is a candidate in the full internal state, not a solution of the original field.

Restoring the coupling means evaluating the original equations and letting both energy channels evolve. The initial state and period are then corrected until that original motion repeats. In a computation with existing solutions, a previously corrected passive cycle is a better nearby guess than a newly constructed auxiliary oscillation. The calculations below use such inherited solutions where available. This preserves the distinction between the mechanical construction that organizes a search and the particular predictor used in an executed numerical solve.

The sequence is therefore

$$\begin{aligned} &\text{auxiliary oscillation or known passive cycle} \\ &\quad \longrightarrow \text{candidate initial state and period} \\ &\quad \longrightarrow \text{correction under the original internal dynamics} \\ &\quad \longrightarrow \text{independent periodicity and pose checks.} \end{aligned} \tag{4.21}$$

No joint torque is applied during this calculation. Correcting the numerical description of an unforced trajectory is not controlling a robot to force it through a prescribed endpoint.

4.3 A periodicity problem with independent conditions

4.3.1 Fix phase and select one member

Let $\Phi_T(z_0)$ be the flow of Equation (4.14). The physical requirement is

$$\Phi_T(z_0) - z_0 = 0. \tag{4.22}$$

Both z_0 and T are unknown. A time shift would describe the same orbit with another initial state, so a phase condition removes this arbitrariness. For a known nonconstant reference cycle one can use

$$(z_0 - z_{\mathrm{ref}})^{\mathsf{T}} f(z_{\mathrm{ref}}) = 0. \tag{4.23}$$

Alternatively, a transverse Poincaré section determines the event time. The phase condition must be meaningful: the tangent expression above cannot use an equilibrium where $f(z_{\mathrm{ref}}) = 0$.

A further condition selects the desired cycle, for example its energy or its mean longitudinal velocity

$$\bar{v} = \frac{1}{T} \int_0^T v(t)\ \mathrm{d}t. \tag{4.24}$$

This is a body-frame quantity. It is not generally the world-frame translation rate $\Delta x/T$, because heading changes during the motion. Nor is a chosen speed a naturality criterion: straight motion and many actuated trajectories can share the same mean velocity.

Shooting evaluates the endpoint by integrating the differential equations; collocation instead approximates a complete trajectory and enforces the equations and boundary conditions together. Both require a phase convention and a family label. Continuation then uses a corrected solution as a predictor for the next one (Peeters et al., 2009; Kuznetsov, 2004).

4.3.2 Use the POE theorem without adding a redundant constraint

On a same-direction joint section, the theorem replaces the joint-rate difference by I_{POE}. The two body velocities must also repeat. These three section conditions are not all independent at a conservative periodic orbit: the first integral of total energy imposes a relation among them. Adding an energy condition to every state condition as if it were independent would give an incorrect count of the numerical problem.

The scalar implementation makes that dependence explicit. For a prescribed nonzero $\bar{v}$, set $h = 1/|\bar{v}|$ and $s = \mathrm{sgn}(\bar{v})$, and use

$$Y = (\delta, \zeta, u, Q), \qquad \sigma = s\zeta, \quad v = su/h, \quad \omega = hQ/u. \tag{4.25}$$

At the initial event $\delta = 0$, $\zeta, u > 0$; logarithms of ζ and u keep these signs during correction. The exact chart and time-orientation convention are derived in Appendix C. They do not change the physical theorem of the previous chapter.

Here $u = sv/|\bar{v}|$ is forward speed normalized with the selected mean speed and oriented to its sign, while $Q = sv\omega$ combines forward and yaw velocities. The bracket $\langle u \rangle$ below is the time average over one complete cycle; setting it to one sets the prescribed $\bar{v}$. The sign s describes the selected speed sheet, whereas the physical theorem's crossing-direction sign χ is a separate choice.

For the 29-root realization, the four-component residual is

$$\mathcal{R}_{\mathrm{POE}} = \begin{bmatrix} I_{\mathrm{POE}}/e_{\mathrm{sc}} \\ (u^+ - u_0)/u_{\mathrm{sc}} \\ (Q^+ - Q_0)/Q_{\mathrm{sc}} \\ \langle u \rangle - 1 \end{bmatrix} = 0. \tag{4.26}$$

The first three conditions express internal periodicity on the section; the fourth sets the requested mean speed. The positive factors $e_{\mathrm{sc}}, u_{\mathrm{sc}}, Q_{\mathrm{sc}}$ normalize quantities of different magnitudes for the numerical solve; they are not new mechanical parameters. For these 29 roots the implemented scales are $e_{\mathrm{sc}} = 1$, $u_{\mathrm{sc}} = \max(1, u_0)$ and $Q_{\mathrm{sc}} = \max(1, |Q_0|)$; they are frozen when the local Jacobian is evaluated. Three initial coordinates are unknown. Their overdetermined least-squares solve is compatible because one differential relation among the first three conditions follows from energy conservation. The independent reintegrations described below check the full internal state, not only the conditions supplied to the corrector.

For the connected positive-speed calculation one may omit the u condition locally. After $I_{\mathrm{POE}} = 0$ has recovered ζ and Q has repeated, total energy recovers u if its derivative with respect to u is nonzero on the section. At fixed $(\zeta, Q, \bar{v})$ the derivative is

$$\frac{\partial E}{\partial u} = \frac{u}{h^2} - \frac{B_{11}(0)h^2Q^2}{u^3} - \frac{sB_{12}(0)hQ\zeta}{u^2}. \tag{4.27}$$

The implicit-function theorem then makes u locally unique at fixed total energy. This is an exact reason to omit that condition; the derivative is checked to remain separated from zero at the reported nodes, also relative to the magnitudes of its potentially cancelling terms. Full-state reintegration is retained regardless.

4.3.3 Follow neighbouring cycles

The connected calculation solves the three independent equations

$$R(X; \bar{v}) = \begin{bmatrix} [E_\perp(T) - E_\perp(0)]/e_{\mathrm{sc}} \\ (Q^+ - Q_0)/Q_{\mathrm{sc}} \\ \langle u \rangle - 1 \end{bmatrix} = 0, \qquad X = (\log \zeta_0, \log u_0, Q_0/Q_{\mathrm{sc}}). \tag{4.28}$$

This connected calculation uses fixed scales $e_{\mathrm{sc}} = 1$ and $Q_{\mathrm{sc}} = 8$, rather than the root-dependent yaw scale of the preceding 29-root audit. The event supplies T. If $D_X R$ is nonsingular, the branch tangent is

$$\frac{\mathrm{d}X}{\mathrm{d}\bar{v}} = -(D_X R)^{-1} D_{\bar{v}} R. \tag{4.29}$$

It predicts the next initial state; Newton correction then recovers the neighbouring cycle. This predecessor-based construction provides an actual numerical connection, not only a visually ordered collection of independently solved targets. The map derivatives account for the change of event time. Otherwise a sensitivity computed at a fixed time would not be the derivative of the Poincaré map.

When a chosen label becomes singular, pseudo-arclength continuation offers a different local coordinate:

$$\tau_k^{\mathsf{T}}(\mathcal{X} - \mathcal{X}_k) = \Delta a, \qquad \mathcal{X} = (X, \bar{v}). \tag{4.30}$$

Here τ_k is the preceding branch tangent in the augmented unknowns. This is a standard numerical option, not evidence that the present positive-speed family has a fold.

A related conservative continuation formulation introduces an auxiliary isotropic term $\kappa\eta$:

$$M\dot{\eta} + c + \begin{bmatrix} 0 & 0 & U' \end{bmatrix}^{\mathsf{T}} + \kappa\eta = 0. \tag{4.31}$$

This variable can regularize the numerical equations while preserving the physical periodic solutions. Indeed,

$$\dot{E} = -\kappa\eta^{\mathsf{T}}\eta, \qquad 0 = E(T) - E(0) = -\kappa \int_0^T \eta^{\mathsf{T}}\eta \, \mathrm{d}t. \tag{4.32}$$

For a nonconstant cycle the integral is positive, hence $\kappa = 0$. The auxiliary coordinate therefore vanishes at an exact periodic root. It is not physical damping used to maintain the conservative motion. The 49-node result below uses the independent POE formulation Equation (4.28), not this alternative unfolding as an unreported extra physical mechanism.

4.4 The computed two-segment repertoire

4.4.1 Twenty-nine roots realize the scalar result

Using the reference mechanics in Equation (4.9), correction and independent event integration recover 29 primitive passive cycles on the two nonzero-speed sheets. They repeat their complete internal state, conserve total energy, satisfy cyclic POE balance and have nonidentity reconstructed pose. The largest independently integrated terminal-state defect is 2.85×10^{-14} in the normalized coordinates; the largest scaled residual of the repolished roots is 4.57×10^{-15}.

These small numerical values are interpretable because the physical field, event, positivity margins and refined integration are all checked. They are floating-point observations, not exact-arithmetic error bounds. The roots were recovered from stored predictors rather than discovered without seeds. Their value here is to realize the energy criterion as a working periodicity calculation.

The numerical rank also agrees with the expected conservative dependence. At all 29 roots the three section conditions have rank two. Adding the fixed-speed condition gives a rank-three 4×3 Jacobian, with smallest singular value between 0.021917 and 0.104764. Restoring $\bar{v}$ as an unknown gives a 4×4 Jacobian of rank three, consistent with a local one-dimensional family before phase is restored. The rank pattern is stable under the reported changes of finite-difference step (Appendix C). These samples do not connect the positive and negative sheets through $\bar{v} = 0$.

4.4.2 A connected family with speed as its cycle coordinate

A positive-speed root seeds the predecessor-based continuation. Its 49 corrected nodes span $\bar{v} = 1.13279$ to 25.0 without changing the robot. Each node supplies an entire passive waveform, its

primitive period and its pose increment. Figure 4.2 displays these loops as a repertoire: moving around one loop advances phase; moving across the loops selects a different gait. The same four highlighted cycles appear in all three panels. Read a loop in panel (a), locate its period in (b), then compare its joint motion with the others in (c). The vertical coordinate in (a) is the cycle's mean speed, not its instantaneous forward velocity.

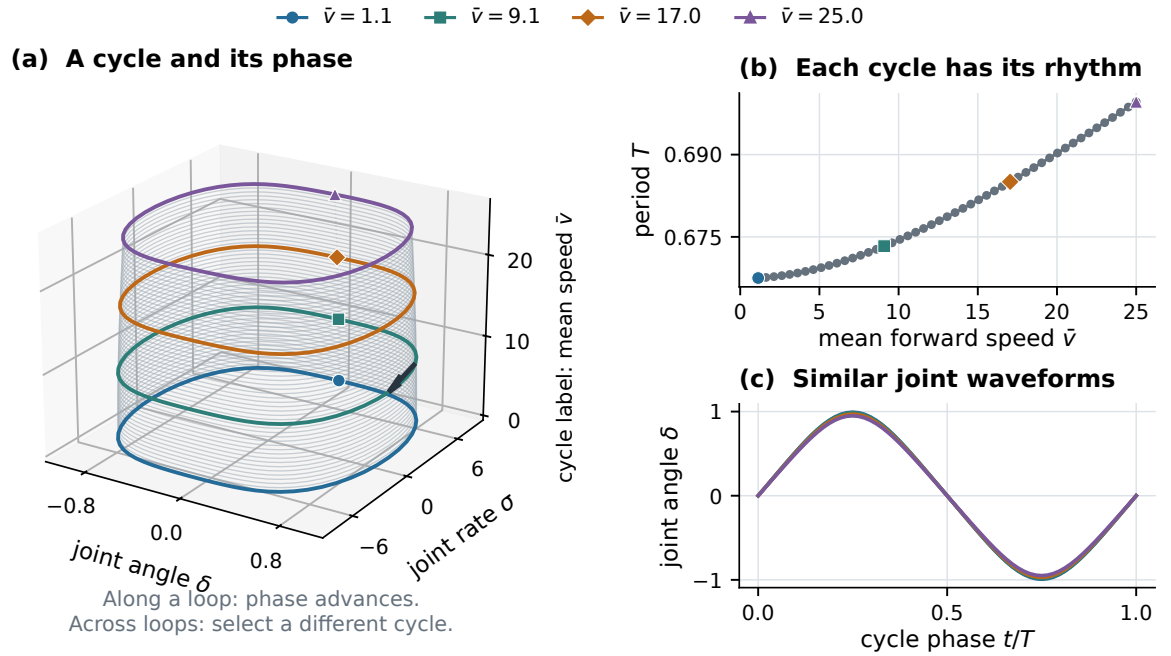


Figure 4.2: A natural-gait repertoire of one 2SEG mechanism. The 49 predecessor-connected cycles differ in mean longitudinal velocity while the masses, geometry, inertia and spring potential remain unchanged. Four cycles are highlighted consistently by colour and marker; their markers in (a) identify the same initial phase. The arrow follows one cycle, while the grey loops show the other computed members. Panel (b) gives every computed period; panel (c) compares four joint waveforms on normalized cycle phase. No interpolated surface or cross-cycle transition is drawn. Velocities and periods use the normalized model scales.

The waveforms remain close in this window even though mean speed varies substantially. Moreover, the larger-speed cycles have slightly longer periods: faster average advance does not require faster joint oscillation. The family therefore organizes more than a rescaling of a prescribed clock; it relates the body's transport to the timing selected by the coupled mechanics.

All nodes satisfy the full-state, primitive-period, total-energy and nonidentity-pose checks. The maximum refined state defect is 6.53×10^{-13}. The smallest singular value of the scaled shooting Jacobian is at least 0.024776, and the nonzero derivative in Equation (4.27) is resolved throughout the sampled set. The branch tangent has a nonvanishing speed component, with minimum normalized magnitude 0.99290. Speed is therefore an evidenced local cycle coordinate for this computed branch, not just a colour assigned to it.

Restoring phase gives

$$z = Z(\bar{v}, \varphi), \qquad 0 \le \varphi < 2\pi. \tag{4.33}$$

The independence of the branch tangent and the flow direction is checked at the sampled cycles. Under the regularity and embeddedness conditions of Appendix G, the sweep is a local NLM. The computation supports this interpretation over its sampled continuation window; it does not provide rigorous interval bounds between every pair of nodes.

This is the physical benefit of the family picture: one mechanism admits a collection of mechanically organized motions at different average speeds. It is not yet a transition law. In the passive

model, time evolution stays on the chosen cycle; changing cycles requires an intervention and may leave the conservative family during the transition. The later resonance and speed-control chapters address maintenance under actuation, not passive jumps between these loops.

4.4.3 What the periodic internal motion does in the plane

The pose equations give

$$\Delta\theta = \int_0^T \omega \, \mathrm{d}t, \tag{4.34a}$$

$$\Delta x = \int_0^T v\cos\theta \, \mathrm{d}t, \qquad \Delta y = \int_0^T v\sin\theta \, \mathrm{d}t. \tag{4.34b}$$

Heading can oscillate while the average displacement accumulates. This explains why a closed internal loop and an open planar trajectory are compatible. It also explains why $\bar{v}$ and $\Delta x/T$ are different observables: the latter includes orientation over the cycle.

A useful comparison is the straight relative equilibrium with the same $\bar{v}$ and duration. With both motions starting from identity pose, define

$$\Delta g_{\mathrm{base}} = \mathrm{Exp}_{\mathrm{SE}(2)}\left(T\widehat{(\bar{v},0,0)}\right), \qquad \Delta g_{\mathrm{inc}} = \Delta g_{\mathrm{base}}^{-1}\Delta g_{\mathrm{cycle}}. \tag{4.35}$$

Its principal logarithm

$$\mathrm{Log}_{\mathrm{SE}(2)}(\Delta g_{\mathrm{inc}})^{\vee} = (\rho_{\mathrm{inc}}, \phi_{\mathrm{inc}}) \tag{4.36}$$

separates residual translation and rotation after this matched composition. The group operation remains meaningful in the presence of rotation, unlike subtracting a straight world-coordinate distance without a frame convention.

For the 29 corrected roots, the normalized translation-log residual $\|\rho_{\mathrm{inc}}\|/(|\bar{v}|T)$ ranges from 0.00377 to 0.05375, or about 0.38–5.37%, in the reported nonzero-speed set. This describes path modulation relative to a specified straight reference. It is not an efficiency measure or a unique causal fraction of propulsion attributable to the joint. The central result remains dynamical: these distinct moving cycles solve the same passive field.

4.5 Changing the mechanics changes the available oscillations

Finding cycles for one robot and selecting the robot are different design tasks. The first leaves its parameters fixed. The second asks whether a change of inertia, potential or geometry makes a desired organization easier to obtain. The same periodicity calculation can evaluate the motions produced by each design, without assuming their waveforms in advance.

4.5.1 An exact inertial relation

Consider the restricted class

$$\ell_1 = \varepsilon\ell, \quad \ell_2 = (1-\varepsilon)\ell, \qquad m_1 = \varepsilon M, \quad m_2 = (1-\varepsilon)M, \qquad I_i = \gamma m_i \ell_i^2. \tag{4.37}$$

The split ε fixes relative length and mass; γ scales inertia while the spring coefficients remain fixed. This isolates an interpretable mechanical effect rather than varying all parameters.

At aligned shape, the yaw–joint inertia block is $B = \left[\begin{smallmatrix} B_{11} & B_{12} \\ B_{12} & B_{22} \end{smallmatrix}\right]$. Set

$$\begin{gathered} D_B = \det B, \qquad \mathfrak{c} = \varepsilon + (1-\varepsilon)^2, \qquad \mathfrak{d} = (1-\varepsilon)^2, \\ C_2 = B_{22}\mathfrak{c} - B_{12}\mathfrak{d}, \qquad N = B_{11}\mathfrak{d} - B_{12}\mathfrak{c}. \end{gathered} \tag{4.38}$$

The mechanical reason for examining N is

$$B^{-1}\begin{bmatrix}\mathfrak{c}\\ \mathfrak{d}\end{bmatrix} = \frac{1}{D_B}\begin{bmatrix}C_2\\ N\end{bmatrix}. \tag{4.39}$$

The original equations contain the constraint-induced velocity coupling; inertia maps it into angular accelerations. The coefficient N/D_B is its joint-direction component. It is not POE power or exchanged energy, and $N = 0$ is a relation among mechanical parameters, not a joint-angle event.

Writing $\beta = 1 - \varepsilon$, the normalized aligned coefficients are

$$\begin{aligned} B_{11} &= \gamma\varepsilon^3 + (\gamma+1)\beta^3 + \varepsilon^2 + 2\varepsilon\beta^2,\\ B_{12} &= (\gamma+1)\beta^3 + \varepsilon\beta^2, \qquad B_{22} = (\gamma+1)\beta^3. \end{aligned} \tag{4.40}$$

Substitution yields

$$N(\varepsilon,\gamma) = \varepsilon(1-\varepsilon)^2\{\gamma(\varepsilon^2+\varepsilon-1) + \varepsilon(\varepsilon-1)\}, \tag{4.41}$$

$$\gamma_\star(\varepsilon) = \frac{\varepsilon(1-\varepsilon)}{\varepsilon+\varepsilon^2-1}. \tag{4.42}$$

For $0 < \varepsilon < 1$, the positive branch of this locus has $(\sqrt{5}-1)/2 < \varepsilon < 1$. This identifies where the equation $N = 0$ has positive γ, not a restriction on all admissible robot designs.

The linearized transverse equations about the straight state $(0, s_0, 0, 0)$ make its meaning more explicit. At $N = 0$ their characteristic polynomial factorizes:

$$p_{s_0}(\lambda) = \left(\lambda + \frac{s_0 C_2}{D_B}\right)\left(\lambda^2 + \frac{k_2 B_{11}}{D_B}\right), \qquad T_\star = 2\pi\sqrt{\frac{D_B}{k_2 B_{11}}}. \tag{4.43}$$

The second factor is an oscillatory pair whose period is independent of the reference translation speed. The first describes the remaining transverse direction. Appendix B gives the short linearization and determinant calculation. This exact result provides a mechanical prediction to compare with finite-amplitude gaits.

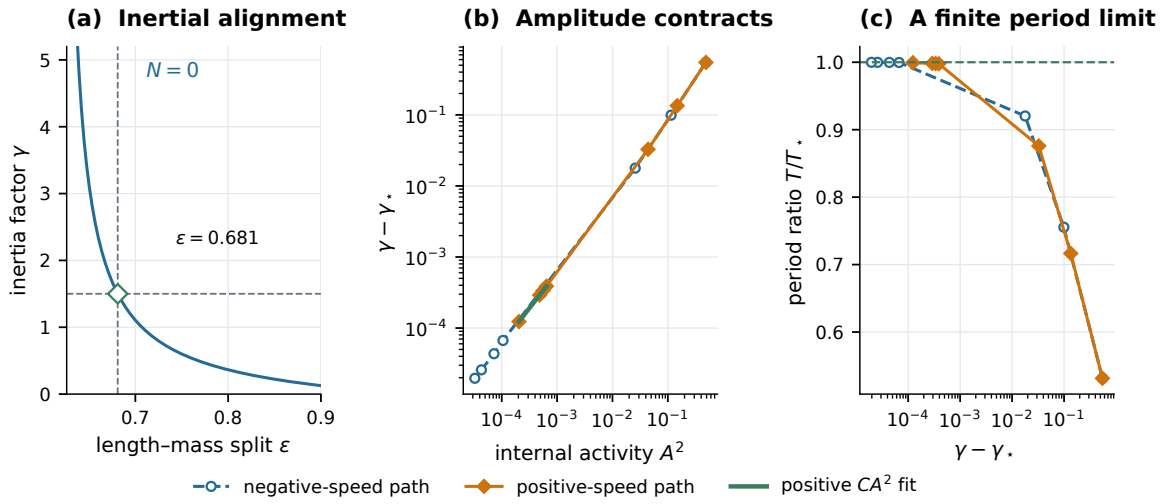
4.5.2 Finite-amplitude cycles approaching the relation

At $\varepsilon = 0.681$, $k_2 = 1$ and $k_4 = 10$,

$$\gamma_\star = 1.5006735, \qquad T_\star = 1.5596234. \tag{4.44}$$

Seven cycles on each of two separately seeded, opposite-speed paths were corrected with the exact passive field. In each path, γ approaches $\gamma_\star$ from above. These are different cycles on different robots: decreasing amplitude across the samples is not a decay occurring during one passive trajectory.

Figure 4.3 connects the analytical prediction to these finite-amplitude motions. Panel (a) locates the relation $N = 0$ in design space; panels (b) and (c) follow the corrected cycles as that relation is approached. With $A = \max_{0\le t\le T}|\delta(t)|$, smaller A^2 accompanies a smaller inertia offset, while $T/T_\star$ approaches one. The circle and diamond markers keep the two opposite-speed paths distinct.



Negative-speed fit: leave-one-out slope span 8.82%, above the 5% criterion.

Figure 4.3: Changing inertia rather than changing phase or gait on one robot. The exact relation $N = 0$ predicts a small-oscillation period $T_\star$. Fourteen corrected passive cycles, seven on each of two opposite-speed parameter paths, approach the relation from above; their joint amplitudes decrease and their periods approach the prediction. The plotted samples do not determine an existence boundary on both sides of the relation.

Both paths show smaller oscillations and periods approaching $T_\star$. On the four near-limit positive-speed samples, the relation

$$\gamma - \gamma_\star \simeq CA^2, \qquad C = 0.6034, \qquad R^2 = 0.999997 \tag{4.45}$$

is numerically well resolved. The opposite-speed path shows the same qualitative trend, but its fitted coefficient is sensitive to the selected samples. The evidence therefore supports the positive-path local fit, not a single robust fitted law shared by both paths.

This calculation illustrates how analysis and periodicity correction complement each other. The inertial relation identifies a place to look and predicts the limiting clock; the nonlinear solve determines actual finite-amplitude motions. No cycle at or below $\gamma_\star$ was tested. Thus the observations do not establish a two-sided bifurcation or a global existence boundary. They do provide a mechanical explanation of how changing inertia affects the natural oscillations that the coupled robot can sustain.

A repertoire and a design question

The output is a set of passive motions with their own timing, not a collection of prescribed joint trajectories. The POE criterion contributes an exact scalar condition; correction enforces the remaining internal periodicity; continuation reveals neighbouring gaits; reconstruction translates their internal motion into a path. On the connected 2SEG family, mean speed and phase provide an accessible local map of this repertoire.

Changing inertia asks a different question and produces a different family of mechanical systems. The exact coupling relation and finite-amplitude results show how the gait calculation can inform that design question. With another joint, several organizations become possible and one scalar energy balance no longer fixes the internal endpoint. The next chapter uses that extra freedom to study two coordinated ways for the same robot to move.

Chapter 5

Two joints, several natural coordinations

A second articulation changes more than the dimension of the equations. It gives the robot several ways to coordinate its internal oscillation: the joints can bend predominantly together, or predominantly against one another. Can these recognizable choreographies persist under the passive mechanics while the same robot advances? As for the double pendulum, equal energy does not imply equal motion. As for the two-segment locomotor, a repeated internal motion need not restore the original pose.

This chapter brings these observations together. We construct the two-coordinate effective oscillator and show why the scalar POE condition no longer characterizes internal periodicity. We then compare in-phase and anti-phase locomotion, before asking how changing the mechanics affects the motions we can recover. The numerical result is local: two organizations coexist at one shared design, with individually verified cycles and local-family evidence. This is already useful for design, without requiring a classification of every possible gait.

5.1 What adding an articulation changes

The three-segment mechanism extends the planar robot of the preceding chapters. A proximal body carries the ideal no-side-slip constraint; two distal segments are connected in series by elastic joints. Here the velocity constraint and elastic potential describe this application, rather than limiting the general motivation for natural locomotion.

The relative joint angles form the shape vector $\boldsymbol{\delta} = (\delta_1, \delta_2)^\mathsf{T}$. The second angle is relative to the middle segment, so the absolute orientations are θ, $\theta + \delta_1$ and $\theta + \delta_1 + \delta_2$. With body velocities $c = (v, \omega)^\mathsf{T}$ and joint rates $w = (\sigma_1, \sigma_2)^\mathsf{T} = \dot{\boldsymbol{\delta}}$, the internal state is

$$z = (\delta_1, \delta_2, v, \omega, \sigma_1, \sigma_2)^\mathsf{T}. \tag{5.1}$$

The pose $g = (x, y, \theta) \in \mathrm{SE}(2)$ is absent from these six coordinates. An internal solution determines its evolution through

$$\dot{x} = v\cos\theta, \qquad \dot{y} = v\sin\theta, \qquad \dot{\theta} = \omega. \tag{5.2}$$

Body velocities therefore belong to the internal state, even though they also reconstruct motion in the plane.

Writing $\xi = (c, w)$, the kinetic energy is $\mathcal{T} = \frac{1}{2}\xi^\mathsf{T} M(\boldsymbol{\delta})\xi$. The matrix follows from the translational and rotational kinetic energies of the three bodies. The joint potential is

$$U(\boldsymbol{\delta}) = \sum_{i=1}^{2}\left(\frac{1}{2}k_{2,i}\delta_i^2 + \frac{1}{4}k_{4,i}\delta_i^4\right) + \frac{1}{2}k_{12}(\delta_1 - \delta_2)^2. \tag{5.3}$$

Thus each joint has linear and cubic restoring torque, and k_{12} adds elastic coupling. No joint trajectory is imposed. Projecting the Lagrange–d'Alembert equations onto admissible velocities gives

$$\dot{\boldsymbol{\delta}} = w, \qquad M(\boldsymbol{\delta}) \begin{bmatrix} \dot{c} \\ \dot{w} \end{bmatrix} + h(\boldsymbol{\delta}, c, w) + \begin{bmatrix} 0 \\ 0 \\ \nabla U(\boldsymbol{\delta}) \end{bmatrix} = 0. \tag{5.4}$$

The vector h contains the inertial terms and the effect of the moving admissible basis. The exact geometry, mass-matrix definition and four components of h are given in Section D.1. Since the ideal constraint supplies no work,

$$E(z) = \frac{1}{2}\xi^{\mathsf{T}} M(\boldsymbol{\delta})\xi + U(\boldsymbol{\delta}) \tag{5.5}$$

is conserved. Numerical values below use normalized model coordinates, without assigning a hardware scale.

5.2 An effective oscillator with two coordinates

5.2.1 The same energy construction, a richer oscillation

The kinetic-energy construction of Section 3.2 applies unchanged, but the shape block is now two-dimensional. At fixed shape and joint rates, minimize kinetic energy over the two admissible body velocities c. Positive-definite inertia gives a unique minimizer $c_\star$. Partitioning M between c and w, and completing the square, gives

$$c_\star = -M_{cc}^{-1} M_{cw} w, \qquad q_c = c - c_\star, \tag{5.6}$$

$$M_\perp = M_{ww} - M_{cw}^{\mathsf{T}} M_{cc}^{-1} M_{cw}, \qquad E_\perp = \frac{1}{2} w^{\mathsf{T}} M_\perp w + U, \tag{5.7}$$

$$E_{\mathrm{prop}} = \frac{1}{2} q_c^{\mathsf{T}} M_{cc} q_c, \qquad E = E_\perp + E_{\mathrm{prop}}. \tag{5.8}$$

Both channels remain part of the internal dynamics. E_{prop} is the energetic contribution of the deviation from the minimizing body velocity, not all body kinetic energy: $q_c = 0$ can still have $c = c_\star \neq 0$. Consequently $E_\perp$ contains kinetic as well as potential energy. These channels are not a relabelling of the kinetic/potential partition.

The auxiliary oscillator is defined by its own Lagrangian,

$$L_\perp(\boldsymbol{\delta}, w) = \frac{1}{2} w^{\mathsf{T}} M_\perp(\boldsymbol{\delta}) w - U(\boldsymbol{\delta}). \tag{5.9}$$

For $D_i M_\perp = \partial M_\perp / \partial \delta_i$, its equations are

$$\begin{aligned} \dot{\boldsymbol{\delta}} &= w, \\ M_\perp \dot{w} &= \frac{1}{2} \begin{bmatrix} w^{\mathsf{T}}(D_1 M_\perp) w \\ w^{\mathsf{T}}(D_2 M_\perp) w \end{bmatrix} - (\sigma_1 D_1 M_\perp + \sigma_2 D_2 M_\perp) w - \nabla U. \end{aligned} \tag{5.10}$$

The shape derivatives of inertia matter at finite amplitude. They ensure that this field conserves $E_\perp$. Conversely, one conserved scalar would not uniquely specify a two-coordinate field; (5.10) supplies the actual auxiliary dynamics.

5.2.2 Modal directions as references for choreography

Near the straight configuration, let $M_0 = M_\perp(0)$ and $K_0 = D^2 U(0)$. The linear modal problem is

$$K_0 u_j = \Omega_j^2 M_0 u_j, \qquad u_i^{\mathsf{T}} M_0 u_j = \delta_{ij}, \qquad j \in \{\mathrm{IP}, \mathrm{AP}\}. \tag{5.11}$$

For the design used below, one vector has entries with the same sign and the other has entries with opposite signs. They define the local *in-phase* (IP) and *anti-phase* (AP) organizations. These names refer to relative joint coordinates, not necessarily to the distal segments' absolute orientations in the world.

The modes provide initial directions, rather than prescribed finite motions. Nonlinear normal-mode calculations extend this small-amplitude picture to periodic oscillations whose period, waveform and modal mixture can change with energy (Kerschen et al., 2009; Peeters et al., 2009). Here this structural-dynamics tool applies to the auxiliary oscillator. Adjoining $c_\star$ to its periodic motion supplies a candidate for the locomotor, which must then be corrected under the original coupled dynamics. The relation $q_c = 0$ used to form the candidate is neither a physical clamp nor an invariant restriction assumed for the full constrained field.

5.3 Why the scalar energy condition is no longer sufficient

The POE principle still requires the effective oscillator and propulsive channel to recover their initial energy allocation over a natural cycle. Along the conservative internal dynamics,

$$P_{\mathrm{POE}} = \frac{\mathrm{d}E_\perp}{\mathrm{d}t} = -\frac{\mathrm{d}E_{\mathrm{prop}}}{\mathrm{d}t}, \qquad I_{\mathrm{POE}}(T) = \int_0^T P_{\mathrm{POE}}\ \mathrm{d}t = E_\perp(T) - E_\perp(0). \tag{5.12}$$

Ideal constraints organize this redistribution without supplying total energy. When transfer is nonzero during a cycle, energy ceded by one channel must be recovered: the exchange has no net gain or loss. The necessary balance is unchanged. What changes is the information carried by one energy value.

For 2SEG, fixing the joint angle and the sign of its velocity left one positive speed to determine. Effective energy was a strictly increasing quadratic function of that speed. With two joints, an energy measurement cannot distinguish all allocations among joint angles and rates. Figure 5.1 first shows the difference while fixing both joint angles: positive crossing direction picks one scalar rate, but leaves several two-joint rate vectors at the same energy. It complements the section argument in Figure 3.2. On the actual 3SEG Poincaré section, only the first angle is fixed, so the remaining freedom is larger still.

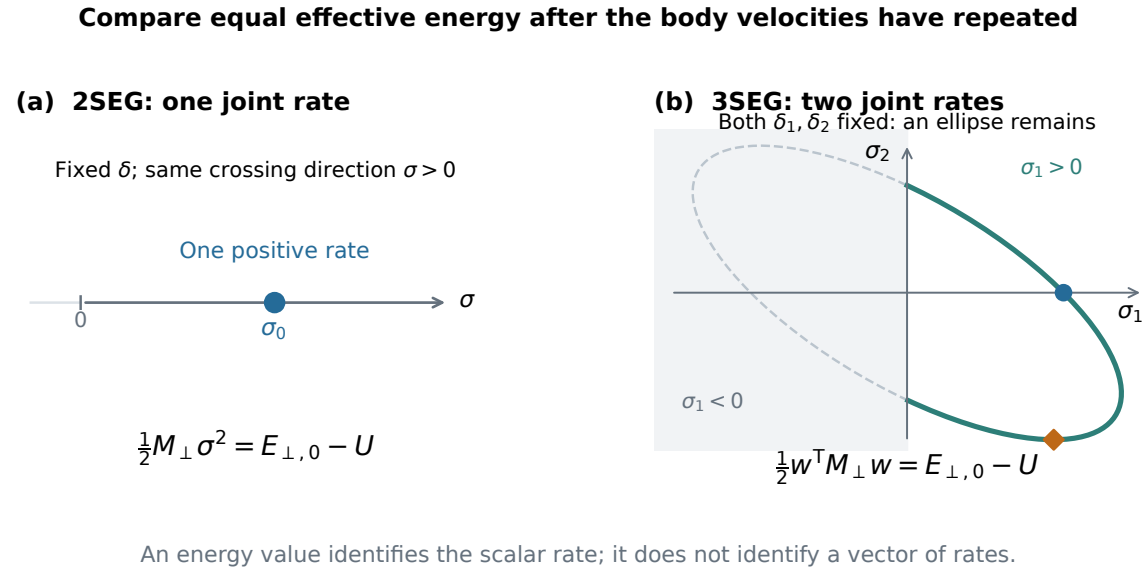


Figure 5.1: Why equal effective energy no longer identifies the joint motion. Assume the body velocities (v, ω) repeat. At fixed shape and positive crossing direction, one effective rate is uniquely identified in 2SEG (left). Even if both 3SEG angles repeat, the positive-definite quadratic energy allows an arc of different rate vectors with $\sigma_1 > 0$ (right); the two markers have equal energy. The ellipse uses the aligned-shape inertia of Equation (D.13), at an illustrative kinetic energy. It is a fixed-shape energy slice, not a trajectory or a NLM, and does not assert that its points are dynamically reachable endpoints.

Proposition 5.1 (The two-joint equal-energy surface)**.** *Consider the oriented section*

$$\Sigma_3^+ = \{\delta_1 = 0,\ \sigma_1 > 0\}. \tag{5.13}$$

After the body velocities have regained their initial values, the remaining endpoint coordinates are $y = (\delta_2, \sigma_1, \sigma_2) \in \mathbb{R}^3$*. If* $M_\perp$ *is positive definite, then, near a point* y_0 *on this section, the equal-effective-energy set*

$$\{y : e(y) = e(y_0)\}, \qquad e(y) = E_\perp(0, \delta_2, \sigma_1, \sigma_2), \tag{5.14}$$

is a two-dimensional surface. Thus $I_{\mathrm{POE}} = 0$ *alone does not identify all three endpoint coordinates; additional relations are needed to infer complete internal periodicity.*

Proof. On the section, $w \neq 0$ because $\sigma_1 > 0$. Consequently $\partial_w E_\perp = M_\perp(\boldsymbol{\delta})w \neq 0$. The gradient of e in its three coordinates is nonzero. The regular-level-set theorem gives dimension $3 - 1 = 2$. □

This states what the energy condition determines, not that the dynamics can reach every point of that surface. A special invariant relation could restrict the reachable endpoints further, but none is assumed here. Even fixing both joint angles illustrates the obstruction: the positive-definite quadratic kinetic energy defines an ellipse in the two joint rates, rather than one rate selected by its sign. The 2SEG theorem therefore cannot simply become a scalar solver for 3SEG. The equal-energy surface here is a set of possible endpoint coordinates; it is not the NLM swept by phase along a family of solved cycles.

5.4 Computing complete internal cycles

Let Φ_p^T denote the flow of the passive six-state model at mechanical parameters p. A natural locomotion candidate satisfies

$$R_z(z_0, T; p) = \Phi_p^T(z_0) - z_0 = 0 \in \mathbb{R}^6, \tag{5.15}$$

with nonconstant internal motion and a nontrivial reconstructed pose. The propulsive motions below have nonzero translation. All joint angles, joint rates and admissible body velocities repeat; the pose itself is not required to repeat.

5.4.1 Fixing phase and selecting a cycle

Autonomy leaves the origin of time arbitrary, while conservation supplies a dependence among periodicity equations. On Σ_3^+, use the five initial coordinates $(\delta_2, v, \omega, \sigma_1, \sigma_2)$ and the period T as six unknowns. The calculation imposes

$$R_{\delta_1} = R_{\delta_2} = R_\omega = R_{\sigma_1} = R_{\sigma_2} = 0, \qquad E(z_0) = E_0. \tag{5.16}$$

The omitted v equality follows locally from total-energy conservation only where

$$\frac{\partial E}{\partial v} = (M(\boldsymbol{\delta})\xi)_v \neq 0. \tag{5.17}$$

After the other five state components repeat, energy then determines v locally and uniquely. This removes a dependent equation; it does not use scalar POE balance to replace independent joint conditions. In every reported sample, $v(T) = v(0)$ is also checked by independent integration of all six equations.

A nonsingular Jacobian of (5.16) with respect to the six unknowns implies a local family as E_0 varies. Because changing energy selects a different conservative orbit rather than a different phase of one orbit, its branch tangent is independent of the flow direction. Restoring phase on a regular embedded portion gives a two-dimensional local Natural Locomotion Manifold. This is the cycle–phase geometry introduced earlier, now inside a six-dimensional internal state space. Numerical rank is assessed at the computed points; the exact conclusion remains conditional on regularity.

5.4.2 Solving the motion before naming it

An auxiliary modal oscillation can initialize the process, but correction in (5.4) is what yields passive locomotion. The present results reuse earlier computed trajectories as predictors, followed by local collocation correction and fresh integration of the full field. This practical reuse does not establish a continuous ancestry from a particular linear mode to every final cycle. We therefore call the movements IP- or AP-*organized*.

Classification follows the periodicity calculation. The modal vectors give reference coordinates

$$Q_j(t) = u_j^{\mathsf{T}} M_0 \boldsymbol{\delta}(t), \qquad \widetilde{P}_j(t) = \frac{u_j^{\mathsf{T}} M_0 w(t)}{\Omega_j}, \tag{5.18}$$

and time-weighted activity

$$A_j = \left[\frac{1}{T} \int_0^T (Q_j^2 + \widetilde{P}_j^2)\, \mathrm{d}t \right]^{1/2}. \tag{5.19}$$

Dominance of the expected activity, together with the sign of joint-angle and joint-rate correlations, identifies the organization. This permits finite-amplitude waveform changes rather than demanding sinusoids with exactly zero or π phase lag. Numerical thresholds and integration settings are specified in Appendix D.

5.5 Two ways for the same robot to advance

5.5.1 One cycle of each choreography

Figure 5.2 compares two corrected primitive cycles near the same mean body velocity. In AP, one joint angle increases while the other predominantly decreases. In IP, their variations predominantly

agree. The shape projections show that these names alone do not specify a motion: amplitudes differ, and the relation between angles need not be a straight line. A shape projection can also retrace itself; it need not enclose an area to represent an internal-state periodic orbit.

The left panels answer a timing question: which joint bends with or against the other during the cycle? The right panels restore physical angle scales. In particular, the nearly retraced AP projection and the curved IP projection show why a modal name is not a complete waveform description. Rates and admissible body velocities remain essential to the internal orbit even when a shape projection nearly retraces a curve.

The motion in the plane gives a complementary intuition. Figure 5.3 compares displacements over the same observation time. Both coordinated oscillations advance the mechanism, but AP executes more cycles during that time. Its printed reference duration contains three primitive periods; it is not the period of the AP trace in Figure 5.2. Read the silhouettes as successive poses, not as independent robots; their vertical offsets only separate the two rows. The drawing supplies the spatial interpretation of the two joint organizations; the seventeen-cycle calculation below supplies the separate periodicity and displacement checks.

5.5.2 Seventeen independently corrected samples

At the shared design $p_\star$ specified below, eight AP and nine IP samples satisfy primitive full internal periodicity, energy conservation, resolved cyclic POE transfer and nonzero reconstructed translation. One further, higher-energy AP target did not converge under the prescribed correction and is excluded from these seventeen successes. Table 5.1 gives the physical range of this result.

The two speed observables are

$$\bar{v} = \frac{1}{T}\int_0^T v(t)\ \mathrm{d}t, \qquad \frac{\Delta x}{T} = \frac{1}{T}\int_0^T v(t)\cos\theta(t)\ \mathrm{d}t, \quad \theta(0) = 0. \tag{5.20}$$

They are close for these nearly straight motions, but not identical. Negative signs reflect the coordinate direction, not failure to locomote.

All seventeen samples have full numerical rank in the phase-fixed six-equation calculation, supporting local families near the individual cycles. Conditioning values are given in the appendix. The samples were corrected independently: a continuous inter-energy path joining successive samples has not been established. The result is coexistence of both organizations and local-family evidence at one mechanism, not a verified connected branch through every displayed energy.

5.5.3 Seeing the cycle–phase geometry

The broader atlas in Figure 5.4 shows how phase and a choice of cycle organize these motions. Moving around one loop changes phase; moving across the stack selects another cycle. Six projections of each sector reveal this geometry without requiring a picture of six dimensions. Intersections in one projection need not be intersections in the full internal state.

Table 5.1: Independently corrected primitive samples at one shared design, in normalized model coordinates. The translation rate includes heading evolution, unlike the mean body velocity.

Quantity	AP	IP
Number of samples	8	9
Total energy	207.070–219.046	187.512–274.084
Primitive period	.187150–.187242	.852475–.854162
Mean body velocity $\bar{v}$	-8.0953–-7.8704	-9.0638–-7.4966
Translation rate $\Delta x/T$	-8.0953–-7.8704	-9.0636–-7.4963

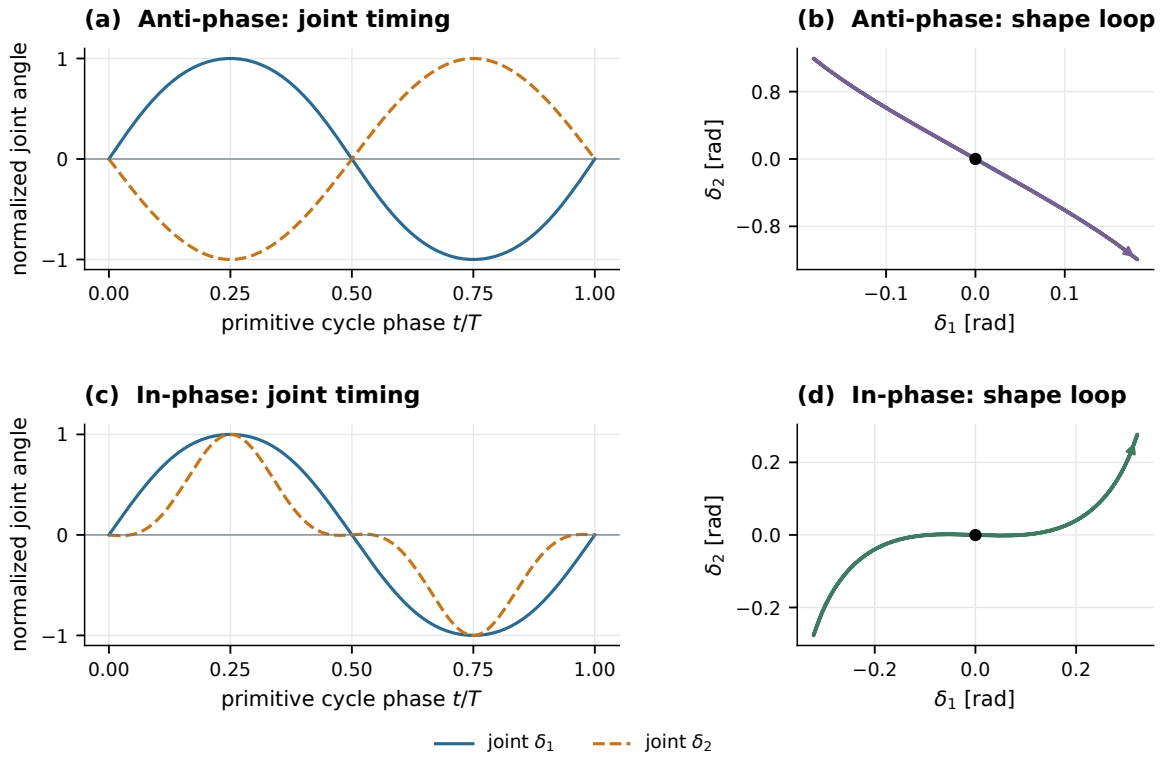


Figure 5.2: One primitive oscillation in each organization: AP above, IP below. Left: joint angles centred at their midrange and divided by their half peak-to-peak excursion; right: the same motion in relative joint angles, in radians, with the initial section marked by a black point and the direction of traversal by an arrow. The periods are $T_{\mathrm{AP}} \simeq 0.187200$ and $T_{\mathrm{IP}} \simeq 0.853107$, with mean body velocities approximately -7.994 and -8.064. These corrected representative trajectories illustrate choreography; normalization does not make their physical amplitudes or periods equal.

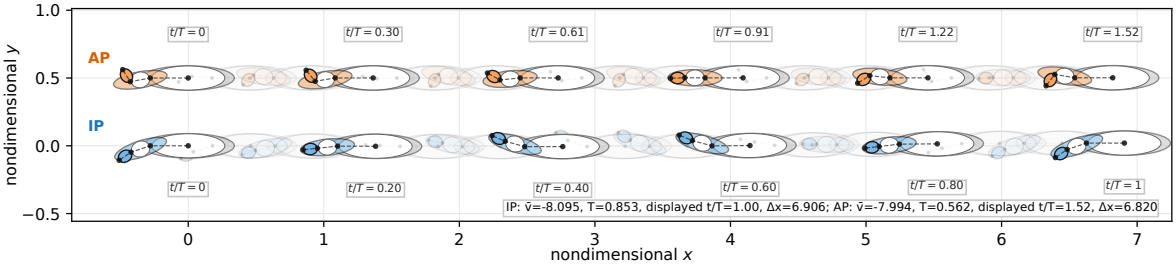


Figure 5.3: Two coordinated motions over a common physical observation interval. AP appears above IP; the rows are vertically offset and faint silhouettes mark intermediate poses. The retained drawing uses $T = 0.561601$ as its AP reference duration, equal to three primitive periods of approximately 0.187200. Its displayed $t/T \simeq 1.52$ therefore represents about 4.56 primitive AP cycles, compared with about one IP cycle. The unequal cycle counts are intentional: the observation times are equal.

Begin with one sector and one highlighted colour. The panels with $\delta_1, \delta_2, \sigma_1, \sigma_2$ as vertical axes show complementary views of the same speed-labelled motion. Next compare the views with $E_\perp$ and signed total energy: they expose channel variation and cycle organization, respectively. Finally compare AP with IP, checking the axis scales as well as the shapes. Their different yaw and joint-rate scales are part of the mechanical contrast, not a plotting inconsistency. The signed-logarithmic colour scale labels mean forward velocity; colour does not indicate phase along a loop.

This atlas has its own supporting sample: 41 AP and 23 IP trajectories, corrected as primitive six-state periodic motions over the broad signed speed range of the original visualization. They are distinct from the seventeen samples above. Each support has an independently integrated endpoint and energy/POE checks. The translucent sheets and highlighted intermediate cuts interpolate the supports; interpolation does not add solved trajectories. Pose and local-rank results for the seventeen-sample study are not automatically extended to the atlas. In particular, 39 atlas supports fail the stricter speed-recovery threshold used to eliminate one periodicity equation in that study, while passing their independent six-state periodicity checks (Section D.4).

Where a speed observable varies regularly along a family, it can select the cycle locally, as developed in Appendix G. Different cycles then provide a repertoire of mean forward speeds. This geometric interpretation does not supply a control law for moving between them: an undisturbed passive trajectory stays on its own cycle.

5.6 Using the computed motions to guide mechanical design

5.6.1 From finding a motion to choosing its mechanics

Until now the mechanics were inputs. For design, the question is reversed: which masses, inertias, geometry and elastic properties make useful coordinated motions available? A visible auxiliary mode is not yet a natural gait. Its energy allocation may be compatible while the coupled body and joint motion still fails to repeat.

The exploratory search first considered individual levers. Proximal inertia was suggested by an energy-exchange diagnostic: adjusting it could remove that diagnostic on a fixed IP candidate without recovering the desired moving IP family. This motivated coordinated changes of joint elasticity, coupling and distal mechanics. The resulting design combines a dominant proximal body with short, light distal segments, asymmetric elasticities and weak inter-joint coupling. These explorations explain the choice of levers; the evidence for the present natural cycles comes from the full dynamical calculations.

Table 5.2 compares that shared design with the reference mechanism. Group its changes into proximal inertia G, joint elasticity/coupling K, and distal mechanics D. Applying just one complete block to the reference mechanism did not recover a nontrivial periodic anchor in six bounded attempts: five ended in straight motion and one did not converge. The combined design supports both organizations. This motivates a local comparison starting from its available cycles.

5.6.2 Changing one block from the shared design

For each parameter in the selected block, define

$$p(\alpha) = (1-\alpha)p_\star + \alpha p_0, \qquad 0 \le \alpha \le 1, \tag{5.21}$$

while other coefficients stay unchanged. Thus $\alpha = .25$ moves one quarter of the way from the shared design toward the reference; it is not a 25% relative change in every coefficient.

At each design, four questions are distinguished: can an auxiliary modal oscillation be calculated; is POE balance obtained; does the complete passive internal state repeat with displacement; and is there local-family evidence at these same mechanics? Partial success then remains informative:

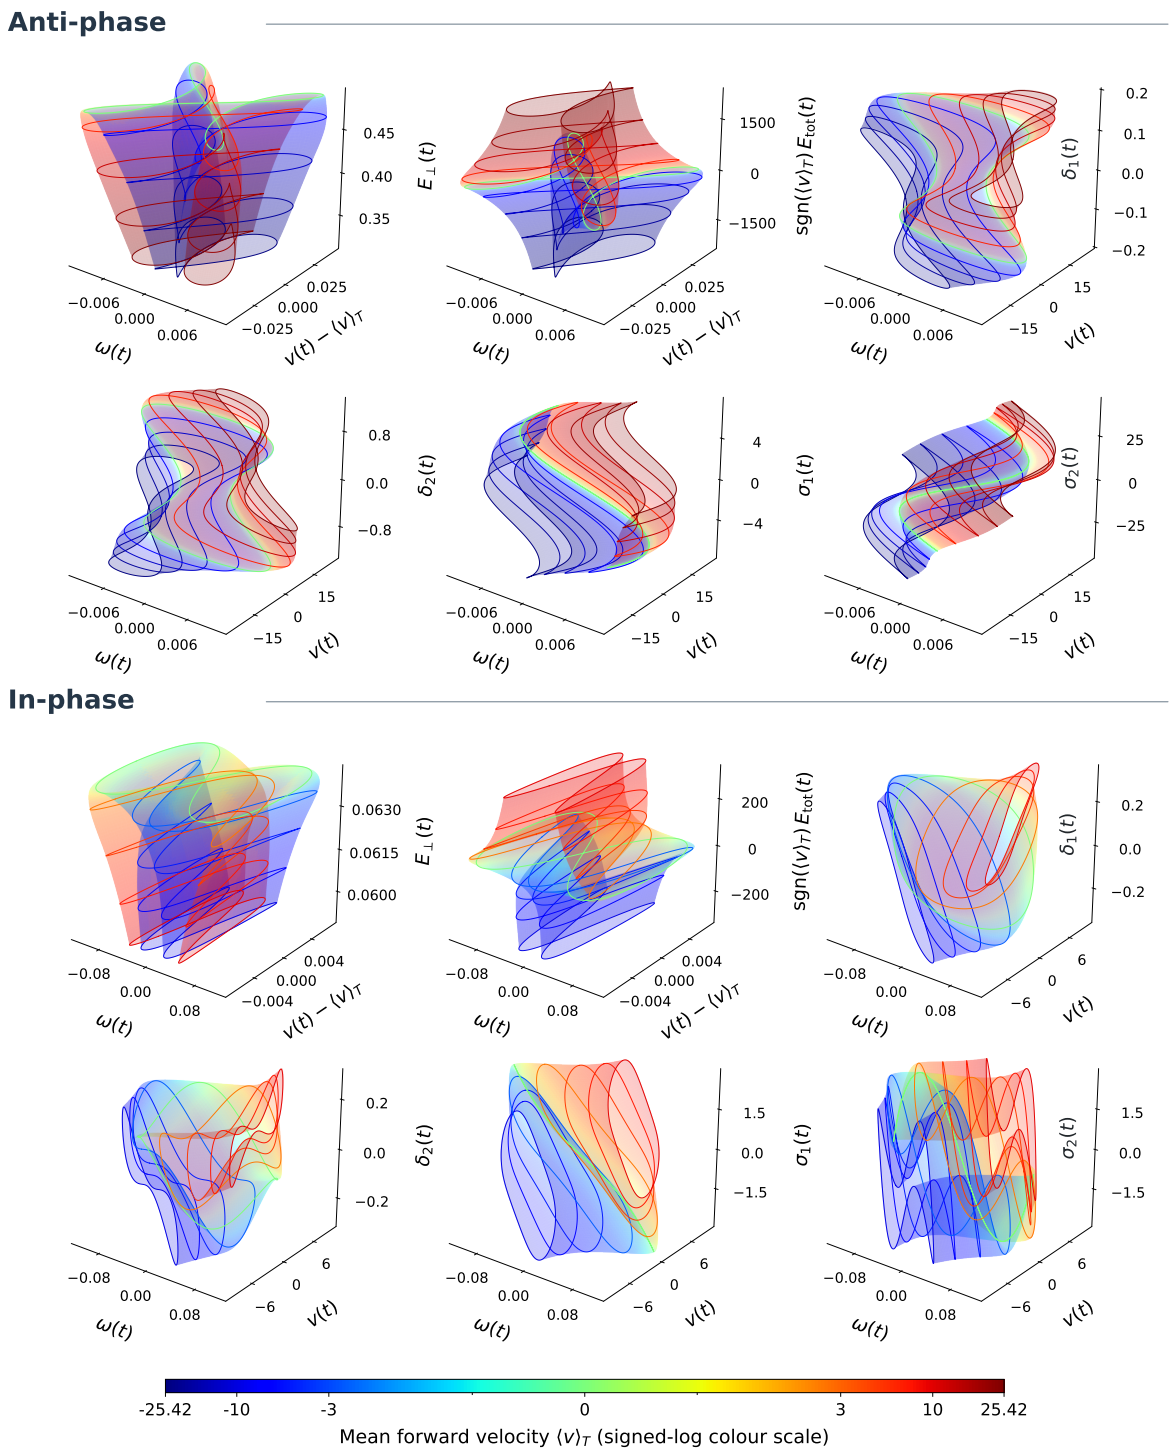


Figure 5.4: A state-space atlas of coordinated cycles. The upper two rows use 41 primitive AP supports, and the lower two rows 23 primitive IP supports, at the same design. Each sector appears in six projections, with nine highlighted cuts at symmetric speed labels. Colour denotes mean body velocity on a signed-logarithmic scale. Energy projections centre $v(t)$ by subtracting its cycle mean; other projections retain $v(t)$. Signed total energy separates velocity directions, not negative physical energy. Sheets and intermediate cuts are visual interpolation. These 64 supports are separate from the eight-AP/nine-IP local-family sample.

Table 5.2: Mechanical choices in the three-segment study. Values are rounded normalized parameters. Blocks distinguish proximal inertia (G), elastic properties (K), and distal mass, inertia and geometry (D), rather than three individual scalar parameters.

Quantity	Reference p_0	Shared $p_\star$	Block
m_1	6	6	fixed
I_1	.3	.418642	G
(m_2, m_3)	$(1.8, 1.2)$	$(.394571, .276554)$	D
(I_2, I_3)	$(.04056, .01936)$	$(.00608070, .000314296)$	D
(l_1, L_2, L_3)	$(.28, .52, .44)$	$(.28, .150076, .105188)$	D except l_1
$(k_{2,1}, k_{2,2})$	$(1, 1)$	$(.652435, .300805)$	K
$(k_{4,1}, k_{4,2})$	$(10, 10)$	$(5.25842, .449143)$	K
k_{12}	.2	.00118731	K

an energy balance alone is not renamed a gait, and a failed correction is not renamed a proof of nonexistence.

The proximal-inertia change retains AP cycles at both displayed points, whereas the first IP calculation retains only an auxiliary oscillation. The elastic change is different: at the quarter point both organizations remain available. At the halfway point neither calculation reaches a complete passive cycle, although IP retains an energy-compatible candidate. Changing the distal block does not recover either organization at the first displayed point. Thus Table 5.3 should be read across each AP/IP pair before comparing the blocks: the same mechanical modification can preserve one organization while the bounded search no longer recovers the other.

These results offer local design guidance. Both organizations tolerate the first elastic adjustment, and the AP search persists farther than the IP search under the chosen inertia reduction. Grouped changes cannot identify a single decisive spring or length. Nor does a stopped numerical path locate a physical termination of a family: the initial guess, representation or corrector may cease to be effective before the motion ceases to exist. Intermediate parameter steps were used where needed to form guesses; a continuous connection across every tabulated design is not asserted.

5.7 What the second joint teaches us

The second joint preserves the energy principle while changing its constructive reach. Cyclic redistribution between the effective oscillatory and propulsive channels remains necessary, but one scalar energy value no longer fixes a vector of joint states. Complete periodicity determines the motion, and modal tools organize the solutions.

At one design the calculations yield both AP and IP natural locomotion, with primitive cycles,

Table 5.3: Motions recovered as one block moves toward its reference values. The entries *Cycle and local family* include the auxiliary, POE, full-periodicity/pose and local-family checks. Other entries describe partial outcomes of the bounded calculation.

Block	α	Anti-phase	In-phase
None	0	Cycle and local family	Cycle and local family
G	.25	Cycle and local family	Auxiliary oscillation only
G	.50	Cycle and local family	Not attempted after prior stop
K	.25	Cycle and local family	Cycle and local family
K	.50	No oscillation or passive cycle recovered	Energy balance only
D	.25	No oscillation or passive cycle recovered	No oscillation or passive cycle recovered
D	.50	Not attempted after prior stop	Not attempted after prior stop

displacement and local regularity assessed numerically. Design comparisons then distinguish an auxiliary oscillation, compatible channel energies and an actual moving periodic solution. The method thus exposes motions that the mechanics can support before a trajectory is prescribed. Two modal directions are useful starting points, not a universal bound on nonlinear gait families: resonances and bifurcations can enrich that organization (Kerschen et al., 2009).

Conservative cycles require no sustained energy input, but their existence does not imply attraction. With dissipation, energy must be supplied if oscillation and advance are to persist. The next chapter turns from finding passive motions to maintaining motions related to them, using actuation timed to accompany the mechanics.

Part III

Dissipation, resonantly maintained locomotion, and local regulation

Chapter 6

Maintaining locomotion by timed energy input

To keep a swing moving, one supplies repeated pushes at the right time rather than prescribing its position at every instant. The swing still determines much of its own motion: the pushes replace lost energy and adapt to its timing. The natural locomotor families suggest the same question for a robot. Having found movements that its mechanics can organize, how can an actuator *maintain* them when energy is dissipated?

The idealization used in the preceding chapters now serves its purpose. Without loss, the effective oscillatory and propulsive channels redistribute energy while their sum remains constant. With positive dissipation, that sum decreases unless energy is supplied. Maintaining a repeated locomotor motion therefore requires work from an actuator, but it need not require a prescribed joint trajectory. Here a sinusoidal joint torque supplies work and timing; the nonlinear dynamics determines the joint waveform, the body velocities and the resulting displacement.

This chapter develops that distinction in two numerical studies. First, a natural 2SEG cycle is continued into maintained responses as loss and actuation are introduced into the same mechanical model. This shows how one cycle deforms. A second study uses a related 2SEG design to explore a broader set of forced responses. It reveals how period, waveform and locomotor output vary, and why a measured phase lag can help an actuator find useful timing. The next chapter turns that observable into feedback.

6.1 Replacing lost energy without prescribing a waveform

The transition can first be understood on a familiar scalar oscillator:

$$m\ddot{x} + c\dot{x} + kx + \alpha x^3 = F\cos(\Omega t). \tag{6.1}$$

Here $m > 0$ is the mass, c the viscous coefficient, and $kx + \alpha x^3$ the restoring force. Multiplying by $\dot{x}$ gives the derivative of $E = m\dot{x}^2/2 + kx^2/2 + \alpha x^4/4$. With $F = 0$ and $c > 0$, the mechanical energy decreases as $\dot{E} = -c\dot{x}^2$ and a non-equilibrium oscillation decays. With forcing,

$$\dot{E} = F\cos(\Omega t)\dot{x} - c\dot{x}^2 = P_{\text{in}} - P_{\text{diss}}. \tag{6.2}$$

If a response returns after the excitation period $T_{\text{exc}} = 2\pi/\Omega$, then its stored energy also returns and

$$\boxed{\int_0^{T_{\text{exc}}} P_{\text{in}}\ \mathrm{d}t = \int_0^{T_{\text{exc}}} P_{\text{diss}}\ \mathrm{d}t.} \tag{6.3}$$

The equality says how much work a maintained oscillation needs over one cycle. It does not say at which times to supply that work or what waveform will result. Those questions require the equation of motion, not energy accounting alone.

The same accounting applies to the locomotor. Ideal constraint reactions still perform no total work: they restrict the admissible velocities and structure the coupling through which the internal oscillator and propulsive channel redistribute energy, but they do not replenish the energy removed by damping. With generalized loss and actuation forces,

$$\dot{E} = \dot{q}^{\mathsf{T}} Q_{\mathrm{loss}} + \dot{q}^{\mathsf{T}} Q_{\mathrm{act}} = -P_{\mathrm{diss}} + P_{\mathrm{in}}. \tag{6.4}$$

The propulsion–oscillator exchange (POE) cancels when the two internal channels are added. It still describes a mechanically meaningful transfer, but it must not be confused with the total input–loss balance: one concerns redistribution *inside* the robot, the other energy entering and leaving it. The conservative identity for each individual channel is not assumed to remain unchanged after external input and loss have been added.

To find the waveform, inspect the internal state once every forcing period, always at the same phase of the torque. This defines a *stroboscopic map* $P_{T_{\mathrm{exc}}}$. A response synchronized one-to-one with the input repeats all its internal variables at these observations:

$$P_{T_{\mathrm{exc}}}(z_0) = z_0. \tag{6.5}$$

For a natural autonomous cycle, shifting the time origin does not change the motion. Here the input supplies a clock: the relation between torque and motion is part of the solution. Finding a fixed point determines a repeated response, while deciding whether nearby motions approach it is a separate stability question.

Read Figure 6.1 in two steps: the linear panel relates timing to the response peak, then the nonlinear panel shows why one forcing frequency need not select a unique motion.

The left panel explains a useful timing relation. A displacement fundamental lagging the force by a quarter-period has its velocity fundamental in phase with that force, so the input supplies positive mean power. This is the intuition behind *quadrature*. The right panel prepares an equally important nonlinear feature: the same input frequency and amplitude may support several responses. Phase, amplitude and performance must therefore be read on the response family, rather than inferred from a single linear curve. Nonlinear response continuation and phase-resonance methods provide established ways to study these features (Peeters et al., 2011; Kuether et al., 2015; Scheel et al., 2018).

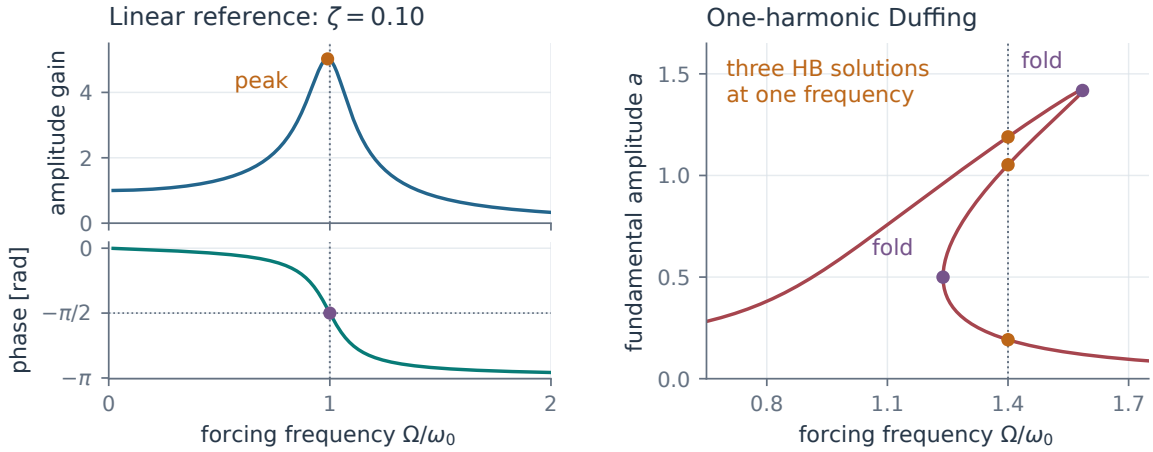
6.2 Following a natural cycle as loss and actuation grow

Start with one of the natural 2SEG cycles already computed. The question is now concrete: can a nearby locomotor response be maintained when small losses are introduced? The geometry, inertias, ideal velocity constraint and quadratic–quartic potential are those of the conservative model in Chapter 4. Only the nonconservative forces change, as shown in Table 6.1. Holding these mechanical ingredients fixed lets us follow the response to added loss and input without also redesigning the robot.

The relation between conservative backbones and weakly forced responses is an established question in nonlinear mechanics (Cenedese and Haller, 2020a,b). Here it is investigated by numerical continuation of the constrained locomotor, not by assuming a persistence theorem for its full family.

Starting from (4.10), the computed bridge ramps viscous loss, isotropic in the nondimensional admissible velocities $\eta = (v, \omega, \sigma)^{\mathsf{T}}$, and synchronized joint torque through the homotopy

$$M(\delta)\dot{\eta} + c(\delta, \eta) + \begin{bmatrix} 0 & 0 & U'(\delta) \end{bmatrix}^{\mathsf{T}} + \lambda c_\star \eta = B_\delta \lambda A_\star \sin\left(\frac{2\pi t}{T_\lambda} + \phi_\lambda\right), \qquad B_\delta = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^{\mathsf{T}}. \tag{6.6}$$



(a) Linear amplitude and output-minus-input phase.

(b) A fold in the first-harmonic Duffing approximation.

Figure 6.1: The one-coordinate reference. In the lightly damped linear case, the response peak lies near a quarter-period displacement lag. At finite amplitude, the first-harmonic Duffing approximation illustrates how a response projection can bend and become multivalued. These panels teach the diagnostic and the fold; they are not predictions for the locomotor models below.

Table 6.1: The mechanics-preserving transition. Geometry, ideal constraint, admissible-velocity basis, inertia, and potential are identical in all three columns; only the generalized nonconservative force changes.

	Conservative	**Damped, unforced**	**Damped and forced**
Generalized nonconservative force	$Q = 0$	$Q = Q_{\text{loss}}$	$Q = Q_{\text{loss}} + Q_{\text{act}}$
Energy implication	$\dot{E} = 0$	$\dot{E} = -P_{\text{diss}}$	$\dot{E} = -P_{\text{diss}} + P_{\text{in}}$

Here $c_\star = 10^{-3}$ and $A_\star = 0.0371395$. Thus $\lambda = 0$ is the original unforced conservative field. For $\lambda > 0$, damping and forcing are ramped together, while the response period T_λ and input phase ϕ_λ are corrected with the orbit. The time origin is fixed by the oriented section $\delta(0) = 0, \sigma(0) > 0$.

The starting point is a numerically refined natural cycle at

$$\bar{v} = 10.43581096, \qquad T_0 = 0.67499472, \qquad \|z(T_0) - z(0)\|_\infty = 2.22 \times 10^{-14}. \tag{6.7}$$

At each positive λ, the initial internal state, period and torque phase are adjusted so that all four internal variables repeat after one input period and the same mean forward speed is maintained. Thus it is the response that adapts: neither the initial joint waveform nor the original period is imposed. Nineteen prescribed values connect $\lambda = 0$ to $\lambda = 1$. Every point satisfies the internal periodicity equations to a maximum residual of 7.11×10^{-10}, with mean-speed error below 1.08×10^{-12}.

Figure 6.2 follows the construction from left to right: the period and excursion change, the oscillation deforms without being prescribed, and input work replaces the loss on each corrected cycle.

For weak loss and input, the maintained waveform is close to the natural one: at $\lambda = 10^{-4}$, their normalized distance is 4.66×10^{-5}. By $\lambda = 1$, the period has decreased by 10.529% to 0.60392282, while the joint excursion has increased by 13.724%. Both input work and dissipated work are about 0.10456319 per cycle, agreeing to a relative difference of 2.60×10^{-11}.

The numerical result is a local route from a natural cycle to maintained locomotion with a deformed waveform. It does not establish that every cycle of the conservative NLM persists in this way. At zero input, the torque phase has no physical meaning, although the natural cycle remains well defined.

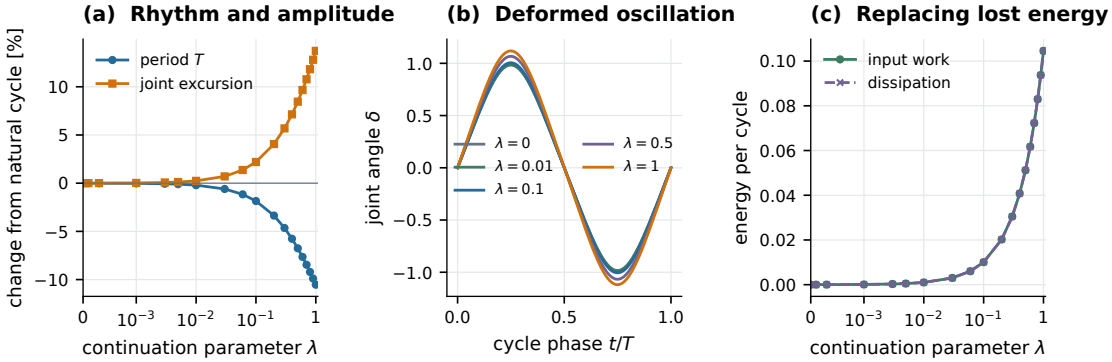


Figure 6.2: One mechanics-preserving conservative-to-forced bridge at fixed mean speed. (a) Period and joint excursion evolve from the conservative anchor along the 19 computed homotopy levels. (b) Section-aligned joint waveforms, plotted against normalized cycle time, approach the anchor as $\lambda \to 0$. (c) Input work replaces viscous loss on each maintained cycle. The internal waveform is an outcome of the corrected dynamics, not a trajectory imposed by the actuator.

Appendix E gives the phase selection, waveform metric, integration checks and complete homotopy specification.

6.3 A broader repertoire of maintained responses

The preceding construction follows one cycle while loss and input grow together. To explore how maintained locomotion varies with excitation, we now consider a second, broader computation: loss is fixed, and torque amplitude and period are varied. Its two-segment kinematics are the same, but its mechanical parameters and potential differ from those of the preceding bridge. The model comes from the forced-response study of Rajaomarosata et al. (2025b). These two studies address complementary questions; they are not consecutive portions of a single continuation.

The configuration and reduced state are

$$q_f = (g, \delta), \qquad g = (x, y, \theta) \in \mathrm{SE}(2), \qquad z_f = (\delta, v, \omega_b, \sigma)^{\mathsf{T}}, \tag{6.8}$$

with carrier-forward velocity v, body-yaw rate $\omega_b = \dot{\theta}$, and joint rate $\sigma = \dot{\delta}$. The same wheel-like velocity constraint removes lateral carrier velocity.

Figure 6.3 separates the kinematic mechanism from its energy accounting. One joint supplies work, whereas dissipation acts in the body and joint velocities; the ideal constraint does not supply the missing energy.

Its projected equations use the verified dissipative sign convention

$$M_f(\delta)\dot{\eta}_f + F_f(\delta, \eta_f) + D_f\eta_f = B_\delta \tau_\delta(t), \qquad \eta_f = \begin{bmatrix} v & \omega_b & \sigma \end{bmatrix}^{\mathsf{T}}, \qquad D_f = c_d I_3, \qquad c_d = 10^{-3}. \tag{6.9}$$

The nondimensional potential is $V_f(\delta) = \delta^2$, so its restoring torque is -2δ. The design vector is $\varepsilon_f = 0.73170844$, $\gamma_f = 0.73680149$. Exact coefficients, scaling, and sign checks are given in Appendix E.

Table 6.2 collects the relation between the models so that the following response plots can be read with a single declared mechanical system in mind.

The imposed torque and its mechanical power are

$$\tau_\delta(t) = A_{\mathrm{exc}} \sin(\Omega t + \varphi_{\mathrm{in}}), \qquad \Omega = \frac{2\pi}{T_{\mathrm{exc}}}, \qquad P_{\mathrm{in}} = \tau_\delta \sigma. \tag{6.10}$$

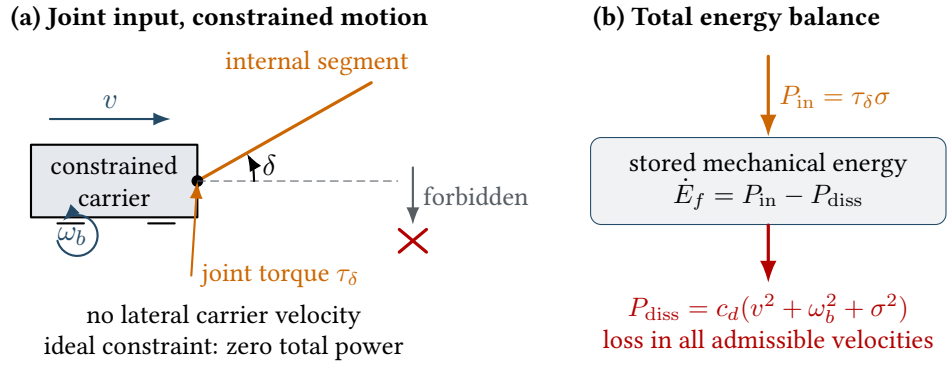


Figure 6.3: Mechanism and energy accounting of the parameter-distinct forced model (conceptual diagram). Joint work maintains total mechanical energy against viscous loss; POE redistributes energy internally and cancels in this total balance. The kinematic pattern matches 2SEG, but the design vector and quadratic elastic law differ from the conservative reference and the bridge in Section 6.2.

Table 6.2: The local bridge and the broad forced-response dataset answer separate questions.

Feature	Same-mechanics bridge	Broad forced-response dataset
Kinematics	Same 2SEG constraint and reconstruction	Same kinematic type
Elastic and inertial design	Exactly the Chapter 4 design	Different design vector and quadratic potential
Forced object	One branch from one conservative anchor at fixed $\bar{v}$	Broad independently computed response set
Question answered	How does one natural cycle change under added loss and input?	Which maintained locomotor responses coexist as excitation varies?

All reported amplitudes are nondimensional. The stored plotting coordinate is $10^3 A_{\text{exc}}$; it is an N mm value only under the optional scaling $k_0 = 1\,\text{N m}$. Viscous dissipation is

$$P_{\text{diss}} = c_d \left(v^2 + \omega_b^2 + \sigma^2\right). \tag{6.11}$$

The actuator acts at the joint, so its instantaneous power is torque times joint rate. Loss also depends on body motion. The mechanical coupling is therefore essential: work supplied through a single articulation supports a response involving all admissible velocities. For a periodic solution these powers satisfy (6.3); the balance follows from the equation of motion and periodicity rather than adding a second independent selection condition.

The pose remains open and is reconstructed from

$$\dot{x} = v\cos\theta, \qquad \dot{y} = v\sin\theta, \qquad \dot{\theta} = \omega_b. \tag{6.12}$$

Thus a forced locomotor response can close in z_f while accumulating a pose increment $\Delta g = g(0)^{-1} g(T_{\text{exc}})$.

6.4 From repeated responses to a locomotion manifold

Suppose the joint angle, joint rate and body velocities have been measured at the beginning of one torque cycle. A maintained synchronized response brings all of them back to those same values after that cycle, while the reconstructed body can have advanced. With z_0 denoting this internal state, the periodicity equation is

$$\mathcal{R}_{\text{per}}(z_0; A_{\text{exc}}, T_{\text{exc}}) := \Phi_{T_{\text{exc}}}^{A_{\text{exc}}, T_{\text{exc}}}(z_0) - z_0 = 0. \tag{6.13}$$

For one amplitude A and period T, this is four equations for four initial state components. Allowing A and T to vary makes it possible to follow neighboring solutions. Numerical continuation follows these periodic responses, including solutions that would not be reached by simply starting an arbitrary trajectory and waiting for transients to decay. An equivalent autonomous forcing-clock formulation is specified in Appendix E.

For the fixed forced model, define the synchronized-response zero set and its transported subset

$$\begin{aligned} \mathcal{S}_R &= \left\{(z_0, A, T) \in \mathbb{R}^4 \times \mathbb{R}_+^2 : \mathcal{R}_{\mathrm{per}}(z_0; A, T) = 0\right\}, \\ \mathcal{M}_R &= \{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}, \qquad G = \mathrm{SE}(2). \end{aligned} \tag{6.14}$$

The first set contains synchronized internal responses; the second retains those whose pose increment is not the identity. In the computed examples below, the increment includes a nonzero translation, so the robot advances rather than merely oscillating in place.

The geometry follows from a simple count made precise by the implicit-function theorem. There are six variables (z_0, A, T) and four periodicity equations. Where their Jacobian has rank four, the solutions form a local two-dimensional surface in this orbit–forcing space. The nonidentity-pose condition is open. A regular embedded local component of the transported set is called a *Resonant Locomotion Manifold* (RLM). A fixed-amplitude slice, when transverse, is the familiar one-dimensional frequency-response curve.

This differs from the conservative NLM geometry. There, one coordinate labels a natural cycle and phase locates the state on that cycle. Here each point represents a maintained response sampled at a fixed input phase, together with its forcing parameters. Sweeping the input phase adds a third dimension in extended state–clock space, not another coordinate of the response surface defined in (6.14).

The adjective *resonant* identifies the primary 1:1 component traced from the neighborhood of the small-amplitude natural period through its nonlinear response ridge. It does not require every response on that component to lie at quadrature or at a speed maximum. Complete continuation Jacobians are not available for the broad archive, so its manifold interpretation remains local and conditional on regularity; the numerical evidence for the responses themselves is given below.

The response population was assembled by continuation in excitation period at fixed amplitude, followed by amplitude continuation from selected responses. At a fold the period reverses direction along the branch, so prescribing a new period can select the wrong neighbor or fail. Pseudo-arclength correction instead fixes a small step along the preceding branch tangent and solves for the period together with the state. It can therefore follow both sides of a fold without treating the fold as an end of the family.

To read these solutions physically, measure their speed, timing, waveform and displacement. With total harmonic distortion (THD) measuring the joint waveform's higher-harmonic content, collect these observables in

$$\begin{aligned} \mathcal{O} : \mathcal{S}_R &\longrightarrow \mathbb{R}_+^2 \times \mathbb{R}^4 \times \mathrm{SE}(2), \\ (z_0, A, T) &\longmapsto (A, T, v_{\mathrm{pk}}, \bar{v}, \phi_{\delta\tau}, \mathrm{THD}_\delta, \Delta g), \end{aligned} \tag{6.15}$$

where v_{pk} is the sampled positive within-cycle maximum, $\bar{v} = T^{-1}\int_0^T v(t)\,\mathrm{d}t$ is the cycle-mean body-frame longitudinal speed, $\phi_{\delta\tau}$ is the deformation phase relative to torque, and THD_δ is the deformation THD. The plots show selected components of this map. Two different responses can therefore appear at the same location in a speed–period plot, just as different points on a surface can overlap in a projection. A fold of a plotted observable does not by itself mean that the underlying response surface is singular.

6.4.1 Phase measured from deformation and torque

The offline phase uses joint deformation as output and applied torque as input. It measures the timing between these two signals, not the locomotor's position along a natural cycle or a phase difference

between two segments. Each adaptive-mesh cycle is first interpolated periodically onto a uniform grid with the duplicated endpoint removed. Centered analytic signals are then used to form the circular mean of their pointwise phase difference. The convention is

$$\phi_{\delta\tau} = \mathrm{phase}_{\mathrm{H}}(\delta) - \mathrm{phase}_{\mathrm{H}}(\tau_\delta) \quad \text{wrapped to } [-\pi, \pi), \qquad \phi^{Q}_{\delta\tau} = -\frac{\pi}{2}. \tag{6.16}$$

Here $\mathrm{phase}_{\mathrm{H}}$ is shorthand for the specified Hilbert analytic-signal statistic, not a pointwise phase on an arbitrary transient. The resampling, circular-mean formula, and harmonic extraction are specified in Appendix E. Chapter 7 uses a different online fundamental estimator with the same signal order and target.

6.5 How excitation changes the locomotor response

At fixed amplitude, the positive peak speed varies with excitation period, and the response curves bend enough that the same input coordinates can correspond to several period-matched responses.

Figure 6.4 displays this projection for the lowest 15 forcing-amplitude levels, $0.00204965 \leq A_{\mathrm{exc}} \leq 0.03551893$. Following one colored curve changes period at fixed torque amplitude; moving between curves changes amplitude. Near the ridge, the speed-maximum circles and quadrature crosses almost coincide. That visual proximity is quantified below, rather than interpreted as equality at every amplitude.

The complete computation contains 13 782 responses at 25 amplitudes over $A_{\mathrm{exc}} = 0.00204965$–$0.07357499$. Their stored internal and clock states agree at the beginning and end of each period. A separate check also propagates every stored interval, including the closing interval, with the declared equations of motion. The largest scaled discrepancy is 1.716×10^{-5}; cycle input and loss agree to a maximum relative discrepancy of 7.583×10^{-6}. These checks support the synchronized response computation at its numerical resolution. The complete amplitude list and verification methods are in Appendix E.

At the straight rest state, eliminating yaw gives the effective shape inertia $M_{\mathrm{lin}} = 0.02582240489$. Since $V_f = \delta^2$, the linear period is

$$T_{\mathrm{lin}} = 2\pi\sqrt{\frac{M_{\mathrm{lin}}}{2}} = 0.7139424640. \tag{6.17}$$

At finite amplitude, the period associated with maximum peak speed shifts away from this reference. The shift is non-monotone: 3.7940% at $A_{\mathrm{exc}} = 0.03551893$, 4.4691% at $A_{\mathrm{exc}} = 0.03969468$, and 3.8114% at $A_{\mathrm{exc}} = 0.07357499$. This change is a property of the complete coupled response, not a correction that can be read from spring stiffness alone.

A fold means that a vertical line in Figure 6.4 can meet one fixed-amplitude curve more than once. These intersections are distinct periodic solutions of the boundary-value problem at the same forcing coordinates. They need not share waveform, mean speed, phase, or pose increment. Their coexistence does not imply that every solution is attracting or reachable from the same initial condition.

6.6 Supplying energy at the right phase

The response diagram tells us which periodic motions exist numerically. It also suggests a signal an actuator could measure while the robot moves: the phase lag between its joint angle and the applied torque. To see why this timing matters, first recall the exact work requirement for the forced model:

$$\int_0^T \tau_\delta \sigma \, \mathrm{d}t = \int_0^T c_d(v^2 + \omega_b^2 + \sigma^2) \, \mathrm{d}t. \tag{6.18}$$

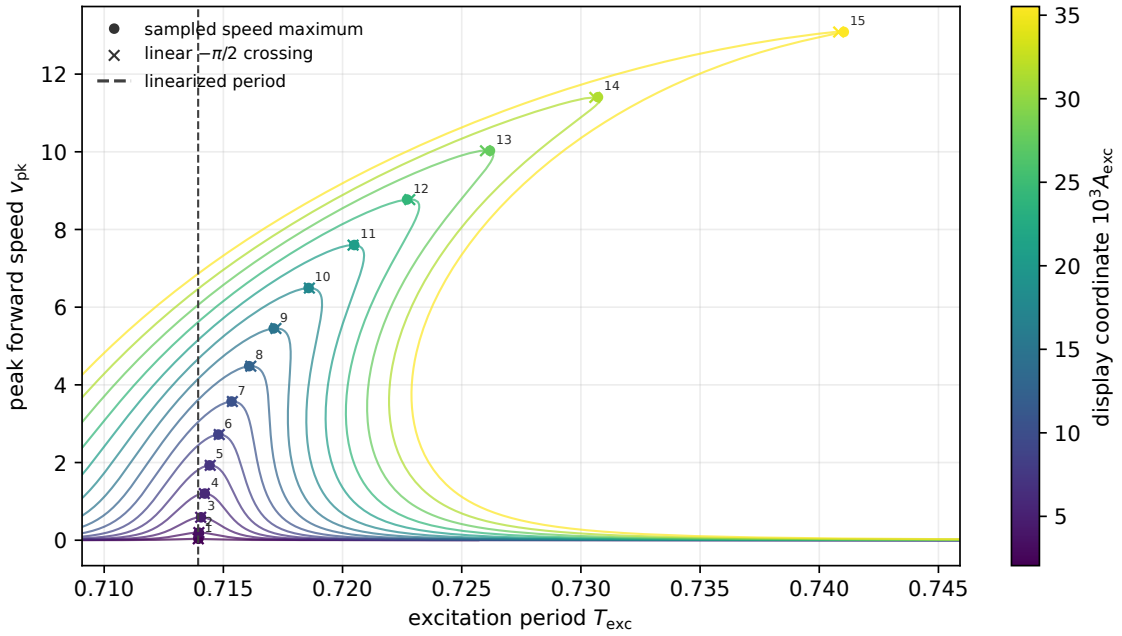


Figure 6.4: Positive peak-forward-speed responses for the lowest 15 forcing-amplitude levels, $0.00204965 \leq A_{\mathrm{exc}} \leq 0.03551893$. Each curve is a continued set at one fixed nondimensional amplitude; the integer beside its peak orders the levels from smallest to largest, so the identification does not depend on color. Circles mark sampled speed maxima, crosses mark piecewise-linearly interpolated $-\pi/2$ phase crossings, and the dashed line is the exact linearized period. Color denotes the display coordinate $10^3 A_{\mathrm{exc}}$.

The actuator supplies work through the joint, while losses occur in all three admissible velocity components. The effective oscillatory and propulsive channels remain coupled: their energy transfer allows joint work to support body motion as well as internal oscillation. The ideal constraint organizes this coupling without adding energy to the total balance.

Now isolate the fundamentals to understand a favorable timing relation. If the torque and joint-angle fundamentals, with positive amplitudes, are

$$\tau_1(t) = \widehat{\tau}\cos(\Omega t), \qquad \delta_1(t) = \widehat{\delta}\cos(\Omega t + \phi_1), \tag{6.19}$$

then differentiation and averaging over $T = 2\pi/\Omega$ give

$$\frac{1}{T}\int_0^T \tau_1 \dot{\delta}_1 \, \mathrm{d}t = -\frac{1}{2}\Omega\widehat{\tau}\widehat{\delta}\sin\phi_1. \tag{6.20}$$

At $\phi_1 = -\pi/2$, this quantity is positive and largest for fixed Ω, $\widehat{\tau}$ and $\widehat{\delta}$. Positive torque is then synchronized with positive joint rate, and negative torque with negative joint rate: both contribute positive work. This is the mathematical counterpart of pushing a swing in the direction in which it is moving. The angle fundamental lags torque by a quarter-period; its rate fundamental is in phase with torque. This argument motivates a phase target, not a requirement that the full nonlinear angle waveform be a delayed copy of the torque. For a purely sinusoidal torque and an exactly periodic response, Fourier orthogonality makes the other joint harmonics contribute zero net input work over a cycle. They still change the waveform and losses. Moreover, the fundamental amplitude changes along the response family: maximizing this fixed-amplitude power expression is not an optimization of locomotor speed or work per distance.

Let $\mathcal{A}_{\mathrm{low}}$ denote the lowest 15 amplitude levels, $0.00204965 \leq A_{\mathrm{exc}} \leq 0.03551893$. At each level, the comparison uses the sampled maximum $(T_{\mathrm{r}}, v_{\mathrm{r}})$ with the piecewise-linearly interpolated $-\pi/2$

crossing (T_Q, v_Q). Every level in $\mathcal{A}_{\text{low}}$ contains such a crossing, and

$$\boxed{\begin{aligned} \max_{A_{\text{exc}} \in \mathcal{A}_{\text{low}}} 100 \frac{|v_Q - v_{\text{r}}|}{v_{\text{r}}} &= 0.0635625\% < 0.064\%, \\ \max_{A_{\text{exc}} \in \mathcal{A}_{\text{low}}} 100 \frac{|T_Q - T_{\text{r}}|}{T_{\text{r}}} &= 0.0289515\%. \end{aligned}} \tag{6.21}$$

The largest phase distance from a sampled speed maximum to quadrature on this range is $0.0452127\,\text{rad}$. These bounds concern one estimator, one sampled domain, and the positive peak-speed observable.

The behavior changes at the next five levels, $0.03969468 \leq A_{\text{exc}} \leq 0.05655617$. They contain a sampled quadrature crossing on a branch whose peak speed is approximately 59–72% below the amplitude-level maximum. At the five upper levels, $0.06052866 \leq A_{\text{exc}} \leq 0.07357499$, the sampled curves do not cross $-\pi/2$.

Figure 6.5 now follows the peak-speed response as amplitude changes. Read the same marker shape across its three panels: circles cover the first 15 levels where the quadrature comparison succeeds, squares the next five where the crossing selects a slower response, and triangles the last five without a sampled crossing.

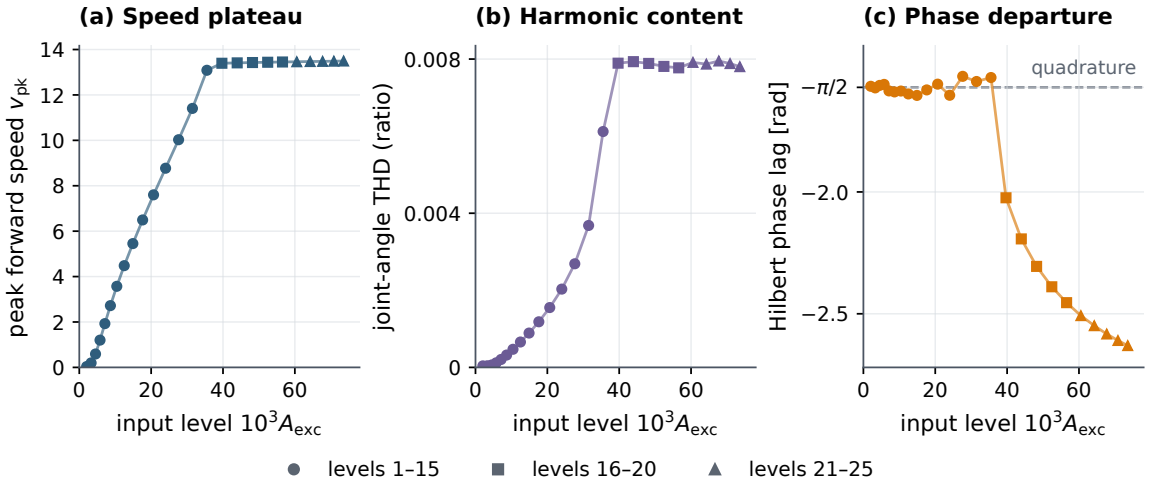


Figure 6.5: Observables at the sampled speed maximum of all 25 forcing-amplitude levels. (a) Peak speed grows and plateaus. (b) Deformation harmonic content changes; THD is a ratio, not a percentage. (c) The output-minus-input Hilbert phase departs from $-\pi/2$ above the first 15 levels. Marker shapes identify the three sampled groups described in the text, not stability classes. All panels use the same rounded nondimensional display value $10^3 A_{\text{exc}}$; connecting lines only guide the eye.

The speed plateau and phase departure are distinct observations: increasing the torque beyond the rising segment scarcely raises peak speed, while the timing of the peak-speed response moves away from quadrature. The changing harmonics also make the actuation philosophy visible. The torque is sinusoidal, but the robot's joint waveform need not be: its nonlinear mechanics determines the shape of the maintained motion. On the lower-amplitude segment, quadrature is a useful local timing marker near the peak-speed ridge. At larger amplitudes, phase alone no longer has that performance meaning.

6.7 A repeated internal motion can keep advancing

The joint and body velocities repeat, but does the robot actually advance? Neither a large within-cycle speed nor phase lock answers that question: an oscillatory velocity can integrate to zero. Integrating

(6.12) gives a nonzero translation for every stored response under both numerical quadrature rules used in the verification. The three reconstructions in Figure E.1, placed beside the reconstruction method in Section E.5, show this spatial result at low, intermediate and high input amplitudes. They confirm the distinction already visible in the natural cycles: the internal motion repeats after each torque period while the carrier advances.

The positive-peak-speed ridge is not automatically an optimum of mean translation, input work per distance, actuator effort, or stability. Nor does phase alone locate a response on a folded map: two responses can have similar phase and different mean speeds. This ambiguity motivates the two coordinates used in Chapter 7: phase adjusts forcing frequency on the fast loop, while mean forward speed adjusts amplitude on the slower loop.

From maintaining a response to choosing its mean speed

The natural-motion construction gave the robot a repertoire. Introducing losses asks how part of that repertoire can be maintained. The first numerical study answers this for one cycle: the motion deforms as torque replaces the dissipated energy. The separate broad response study shows how a related locomotor organizes synchronized motions over forcing amplitude and period, including coexisting responses and changing waveforms.

The actuator therefore has two useful tasks. It can adapt its timing to the response, using the local quadrature relation, and it can change the amount of work supplied to adjust the mean forward speed. The response family helps us understand these choices, but need not be stored as a trajectory for the robot to follow. The next chapter implements the two tasks as feedback loops and tests their ability to maintain locomotion while changing its mean speed.

Chapter 7

Timing the actuation and regulating mean advance

The preceding chapter showed how timed joint torque can maintain a locomotor response without prescribing its waveform. The remaining question is how the actuator should adapt while the robot is moving. Two observations are available: the phase lag between joint angle and torque, and the mean forward speed. The first answers *when should the actuator supply its work?*; the second answers *which locomotor operating point should it maintain?*

We assign one adjustment to each question. A phase-locked loop (PLL) changes the torque frequency to regulate its timing relative to the joint motion. A slower loop changes torque amplitude according to the mean-speed error. The actuator thus accompanies the response: the controller receives neither a prescribed joint waveform nor an online query to the computed response family. Its targets are phase and mean speed, while the robot's dynamics continues to determine the detailed motion.

PLL methods already exploit phase feedback in nonlinear resonance testing and backbone tracking (Peter et al., 2016; Peter and Leine, 2017; Denis et al., 2018). Their relation to control-based continuation has also been studied (Renson et al., 2016; Barton, 2017; Abeloos et al., 2022; Tang et al., 2020). The locomotor-specific development here combines timing regulation with a mean-speed coordinate on the useful response segment identified in Chapter 6. We first explain why both measurements are needed, then give the feedback laws and test their roles in closed-loop simulations.

7.1 Being synchronized is not yet choosing a speed

Consider two maintained joint motions whose fundamental components both lag the torque by almost a quarter-period. Their timing is similar, yet one can have a much larger joint excursion and produce a much higher mean forward speed. The folded response curves of the previous chapter make this ambiguity possible. A PLL can report a small phase error without selecting the intended locomotor response.

The simulations expose this ambiguity directly. An amplitude-staircase test attempted to reach the computed response at $A_{\mathrm{exc}} = 0.0176277$, whose mean speed is 6.2112, while retaining the PLL. Fifteen trials used several low-response starts and target-phase settings. None reached that speed neighborhood: all settled on a low-speed response even though the largest online phase-error root mean square (RMS) was only $0.07138\,\mathrm{rad}$. The loop was controlling a real timing property, but that property did not identify the desired gait.

Let Γ denote a settled forced response. We therefore observe both

$$\Gamma \longmapsto \big(\phi(\Gamma), \bar{v}(\Gamma)\big), \qquad \bar{v}(\Gamma) = \frac{1}{T}\int_0^T v(t)\ \mathrm{d}t. \tag{7.1}$$

Here $v(t)$ is the body's longitudinal velocity in its own frame, and $\bar{v}$ averages that component over a cycle. When heading changes, this is not generally the magnitude of net spatial translation divided by T; pose reconstruction remains the test of net displacement. It is also distinct from the within-cycle peak speed used in the preceding response plots.

The first observable measures timing; the second distinguishes operating points on the selected response segment. To inspect that segment, order its computed responses by increasing forcing amplitude. An *amplitude level* below refers to one of these physical input values, not an index supplied to the online controller. On the first fourteen levels, the pairs

$$(A_{\mathrm{exc}}, \bar{v}) : (0.0020496, 0.021457) \longrightarrow (0.0314870, 10.876395)$$

increase, and every adjacent sampled slope is positive. These data support mean speed as a useful local label: it separates sampled responses whose phases are almost indistinguishable. The next level, $(A_{\mathrm{exc}}, \bar{v}) = (0.0355189, 12.364852)$, supplies the high-speed reference for the numerical control tests.

This is the control counterpart of using mean speed to label natural cycles. A family describes alternative motions, but the passive dynamics does not move spontaneously from one of its cycles to another. Here actuation supplies both the energy needed to maintain a dissipative response and the means to change its operating point. The tests concern the rising response segment just identified, not navigation over every component of the RLM.

7.2 Measuring the phase lag and closing the PLL

The PLL compares the joint-angle signal δ with the applied torque. The latter is generated from an evolving excitation phase:

$$\tau_\delta(t) = A_{\mathrm{exc}} \sin \varphi_{\mathrm{exc}}(t), \qquad \dot{\varphi}_{\mathrm{exc}} = \Omega. \tag{7.2}$$

The joint waveform may contain several harmonics, whereas the phase target concerns its fundamental component. A least-squares fit over approximately one rolling forcing period represents three harmonics of each measured signal. The higher harmonics accommodate waveform distortion; only the fundamental phases enter the feedback error:

$$\phi = \mathrm{wrap}_{[-\pi,\pi)}\big(\phi_\delta - \phi_\tau\big), \qquad e_\phi = \mathrm{wrap}_{[-\pi,\pi)}\left(\phi + \frac{\pi}{2}\right). \tag{7.3}$$

The subtraction is joint-angle phase minus torque phase. Thus a joint-angle fundamental lagging the torque by one quarter-period has $\phi = -\pi/2$ and zero error, as illustrated in Figure 7.1. This is the timing relation whose fundamental joint rate is in phase with the torque. The estimator algebra and precise windows are given in appendix F.

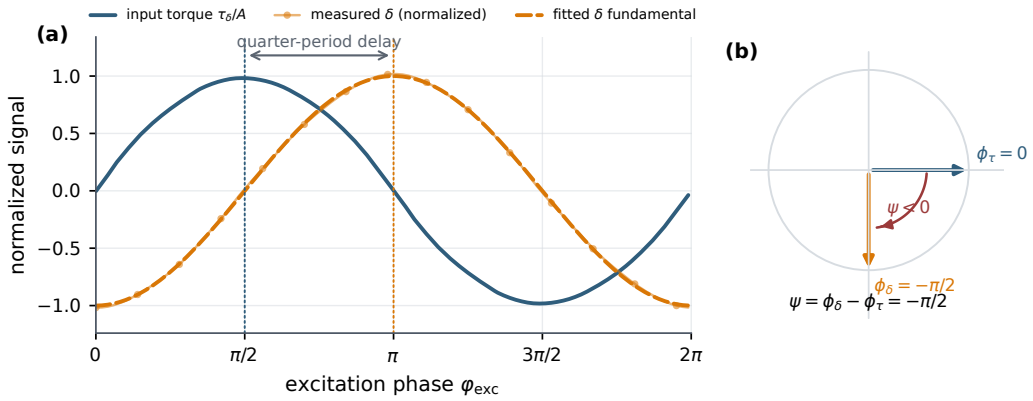


Figure 7.1: Online phase convention. The applied torque provides the excitation clock, and a rolling harmonic fit extracts the fundamental deformation phase. In the output-minus-input convention, the displayed quarter-period delay is $-\pi/2$. Higher fitted harmonics absorb waveform distortion; they are not additional feedback coordinates.

The offline response plots use a Hilbert phase statistic, whereas this online estimator extracts fundamental phases. Their common signal order and quadrature target make their physical interpretations consistent, but not their numerical values identical on distorted or transient signals. The PLL sign is therefore checked using the measurement that actually enters feedback.

A condensed proportional–integral (PI) form of the phase controller is

$$\dot{I}_\phi = e_\phi, \qquad \Omega_{\mathrm{raw}} = \Omega_{\mathrm{free}} + s\big(K_P e_\phi + K_I I_\phi\big), \qquad \Omega = \mathrm{sat}_{[5,11]}(\Omega_{\mathrm{raw}}). \tag{7.4}$$

The proportional term reacts to the present phase error; the integral term adjusts the frequency bias needed to remove a persistent error. The sign s depends on the signal convention and local response slope: the frequency correction must move the measured phase toward its target. Around the amplitude level $A_{\mathrm{exc}} = 0.0176277$, with $\Omega_{\mathrm{exc}} = 8.7439$, a positive e_ϕ must increase Ω, hence $s = +1$. With that sign, initial detunings of -0.5%, $+0.5\%$, and -3% ended with an absolute tail-refitted phase error below $4.2 \times 10^{-3}\,\mathrm{rad}$ and mean frequency close to the reference. Here one harmonic fit is made over the final quarter of each run; its absolute error is not the RMS of the rolling estimates used online. With the opposite sign and the same -0.5% detuning, that refitted error remained $1.2452\,\mathrm{rad}$ and the mean frequency fell to 0.874 of the reference. This calibrates the feedback sign in that local neighborhood; the tested detunings are not a map of its acquisition basin. Section F.3 gives the preliminary test settings and online RMS values; Section F.2 specifies the subsequent nominal phase–speed controller, including startup and saturation.

Figure 7.2 compares the correct and reversed feedback signs under the same local perturbation.

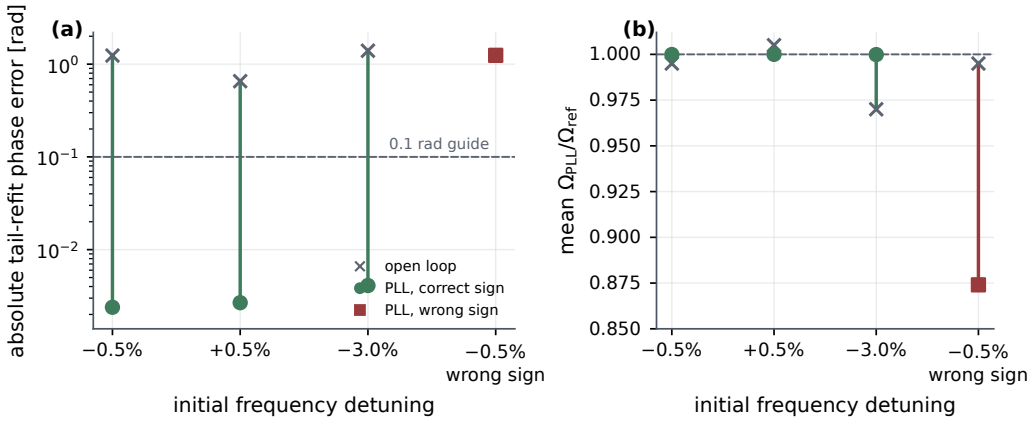


Figure 7.2: Local sign experiment at $A_{\mathrm{exc}} = 0.0176277$. Correct-sign feedback reduces the absolute phase error of a harmonic fit over the final quarter of each run for the three tested frequency detunings and returns the mean frequency near its reference. Reversing the sign under the same -0.5% detuning moves away from that neighborhood.

7.3 Changing mean speed while maintaining phase lock

The two measurements now lead to two coordinated adjustments: Figure 7.3 separates their feedback paths before the laws are written. The blue path adjusts timing; the orange path adjusts mean advance through the error $e_v = \bar{v}_{\mathrm{des}} - \bar{v}$ relative to the requested speed $\bar{v}_{\mathrm{des}}$. A third, visually separate layer belongs to the numerical experiment: it decides whether to keep a completed stage, not how to shape each joint oscillation.

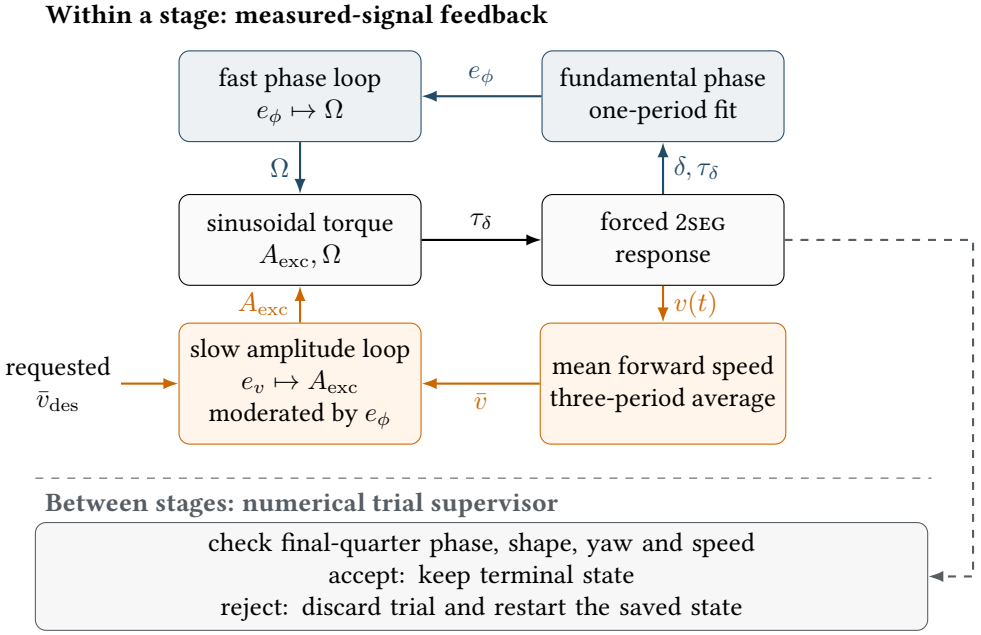


Figure 7.3: Two feedback channels and a separate numerical supervisor (conceptual architecture). Frequency regulates fundamental phase; amplitude regulates mean advance, with phase error moderating its rate. Solid arrows belong to the controller; the dashed route supplies completed-stage data to the simulator's supervisor. Restarting a saved simulated state is not a physical recovery action. No stored waveform or response map is queried online.

The two feedback actions can be summarized as

$$\underbrace{e_\phi \longmapsto \Omega}_{\text{fast phase loop}}, \qquad \underbrace{e_v = \bar{v}_{\mathrm{des}} - \bar{v} \longmapsto A_{\mathrm{exc}}}_{\text{slow speed loop}}. \tag{7.5}$$

Changing Ω corrects the timing of the torque relative to the joint motion. Changing A_{exc} slowly changes which response is maintained. The speed loop estimates mean speed over a rolling three-period window rather than trying to suppress the natural speed variations within each cycle. It moves the operating point while the faster PLL regulates timing.

If desired mean speed exceeds the measured value, the amplitude loop increases torque amplitude; on the sampled rising relation this tends to move the response toward higher speed. An excessively rapid change can, however, disrupt synchronization. The phase error therefore also moderates the rate at which amplitude may change. For the nominal campaign, this weight is

$$g_\phi(e_\phi) = \begin{cases} 1, & |e_\phi| \le 0.20, \\ \dfrac{0.55 - |e_\phi|}{0.35}, & 0.20 < |e_\phi| < 0.55, \\ 0, & |e_\phi| \ge 0.55, \end{cases} \tag{7.6}$$

and the slow law is

$$\begin{aligned} e_v &= \bar{v}_{\mathrm{des}} - \bar{v}, \\ \dot{A}_{\mathrm{exc}} &= \mathrm{sat}_{[-r_A, r_A]}\left(g_\phi K_v e_v - \mathbf{1}_{\{g_\phi = 0\}} \rho (A_{\mathrm{exc}} - A_{\min}) \right), \qquad A_{\mathrm{exc}} \in [A_{\min}, A_{\max}]. \end{aligned} \tag{7.7}$$

The tested values are $K_P = 0.04$, $K_I = 0.08$, $K_v = 2.5 \times 10^{-5}$, $r_A = 2.5 \times 10^{-4}$, $\rho = 10^{-4}$, and $[A_{\min}, A_{\max}] = [10^{-4}, 8 \times 10^{-2}]$. Before a valid phase and three-period speed estimate exist, the amplitude update is held. Here sat clips its argument to the stated interval, and the indicator $\mathbf{1}_{\{g_\phi = 0\}}$ is one only when the phase gate is fully closed. Moderate phase error slows the speed adjustment; larger error removes it and gently reduces amplitude toward $A_{\min}$. These are tuning choices for the tested response segment, not mechanically derived thresholds. The sampled recurrences and the distinction between rolling estimator windows and end-of-stage evaluation windows are stated in Sections F.2 and F.4.

A component ablation at desired mean speed $\bar{v}_{\mathrm{des}} = 6$, reported in Table 7.1, makes the separation concrete.

Table 7.1: Component ablation at the speed-6 command.

Active layers	Achieved $\bar{v}$	Phase-error RMS	Observed role
Phase and speed loops	5.594	8.14×10^{-3}	Moves upward while retaining phase lock
Fixed-amplitude PLL	1.798	7.07×10^{-4}	Locks phase at the starting response
Speed loop without PLL	3.001	0.812	Moves speed but loses the phase marker

With both loops, speed rises while the phase lag remains close to the target. The PLL alone synchronizes the starting low-speed response; amplitude feedback alone changes speed but loses the timing marker. The finite-duration two-loop test reaches 5.594 rather than exactly 6, so it demonstrates the complementary roles of the loops, not zero speed error at every command.

After each finite-duration stage, the stage manager evaluates phase-error RMS, deformation peak-to-peak, mean yaw rate, backward speed progress, and, above the high-speed threshold, proximity to the command. A passing stage retains its terminal plant state, excitation phase, amplitude, and frequency. A failing stage restores those last accepted values and halves the proposed speed increment. Each trial starts fresh estimator histories while retaining the chosen plant and forcing initial state. Thus the demonstrated procedure is a staged simulation protocol, not one uninterrupted implementation of the rolling-window controller. Exact definitions, thresholds, durations, and step-update rules are collected in Section F.4. The tests use measured trajectory quantities only; the stored forcing-amplitude levels are attached after a run for interpretation.

7.4 Closed-loop progression across the rising response segment

The two runs in Figure 7.4 start from rest, supplied with the forcing coordinates $(A_{\mathrm{exc}}, \Omega) = (0.0020496, 8.8006)$, and the stored periodic state at the same amplitude level. Each run accepted 21 stages and reached desired-speed command 12. The terminal mean speeds were 12.068039 and 12.067794, with terminal phase-error RMS 1.119×10^{-3} and 1.106×10^{-3} rad, respectively. The corresponding nondimensional final forcing amplitudes were 0.0364166 and 0.0364088; by mean speed, both endpoints are closest to the stored level at $A_{\mathrm{exc}} = 0.0355189$. The three panels separate the outcome from the protocol: achieved speed approaches the command in (a), the proposed increments become smaller in (b), and the rolling phase error remains small in (c). The smaller high-speed increments are a deliberate stage rule, not a spontaneously measured mechanical threshold.

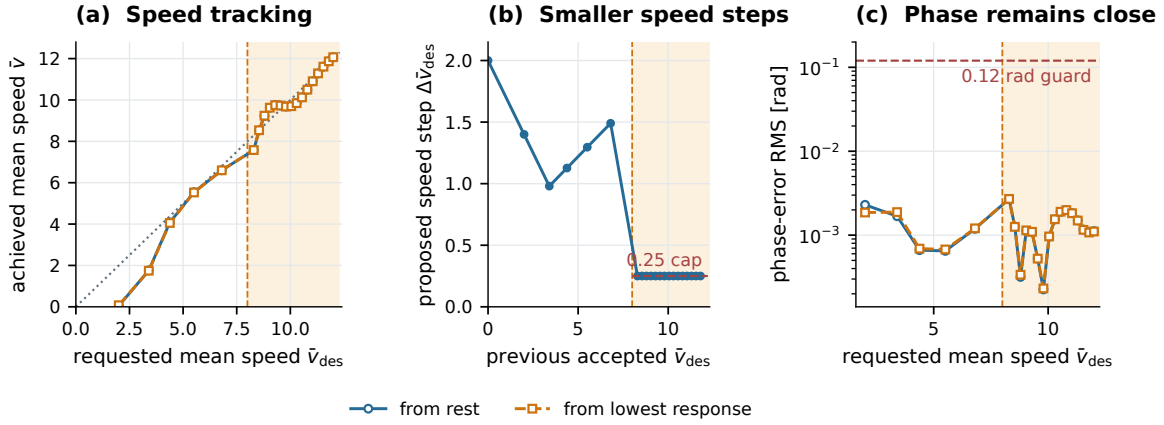


Figure 7.4: Nominal phase–speed progression from rest and from the stored weak-amplitude response at $A_{\mathrm{exc}} = 0.0020496$. Both runs contain 21 accepted stages and terminate near mean speed 12.068. Once the accepted target reaches 8, the proposed increment is capped at 0.25. All phase-error RMS values remain below the acceptance threshold; the closest stored high-amplitude response has $A_{\mathrm{exc}} = 0.0355189$ and is identified only after the run.

Does the controlled endpoint resemble a complete locomotor response, rather than merely matching two observables? For the stored-response start, the last controlled cycle was compared with all 25 stored peak-response cycles at 512 common forcing phases, using every internal component $(\delta, v, \omega, \sigma)$. The nearest profile was the fifteenth amplitude level, with normalized RMS distance 1.7543×10^{-2}; the fourteenth was next at 6.5033×10^{-2}. The last two controlled cycles differed by only 1.1275×10^{-4} in the same normalization.

This comparison supports proximity to the intended high-speed response. It is not an exact periodicity or manifold-membership test: the controller still adjusts its forcing, and only the 25 peak cycles are compared. The metric, endpoint discrepancy and final-stage half-step replay are specified in Section F.5.

All 42 nominal stage trials were accepted, so those two runs do not exercise rejection. A separate stress harness deliberately bypassed the high-speed step cap after the accepted target 8.2935. The direct proposal to 12 was rejected at mean speed 8.9980. The harness then discarded that terminal state, explicitly reused a copy of the last accepted plant and forcing state, halved the proposed increment, and accepted the resulting target 10.1468 at mean speed 8.7457. This tests the simulator, the guard, and one constructed reject–reuse–reduced-proposal sequence; it does not execute the rejection branch of the production stage manager or establish recovery after arbitrary failures. The two nominal endpoints remain the evidence for consistency across the named initializations.

7.5 What the finite-harmonic model tells us

A separate calculation perturbs each sampled response with a constant and $K_{\mathrm{G}} = 5$ harmonics in each of its four internal-state components, giving 44 perturbation coordinates. Projecting the equations onto these coordinates gives homogeneous derivative models for frequency and amplitude feedback. The stored reference has a nonzero projected residual, which is omitted from these linear matrices. The calculation retains the stored reference waveform and approximates the allowed perturbations; it also replaces the rolling estimators by simpler local channels. The resulting closed matrices are displayed in Section F.6. The spectral abscissa is the largest real part among a matrix's eigenvalues; a negative value means asymptotic stability of that finite linear model. This Galerkin order is unrelated to the online estimator order $K_{\mathrm{est}} = 3$.

Table 7.2 reports the resulting spectral-abscissa ranges on the sampled branch segment.

Table 7.2: Spectral-abscissa ranges of the sampled $K_{\mathrm{G}} = 5$ Galerkin controller matrices.

Augmented model	Spectral-abscissa range	Negative samples	Domain
PLL frequency loop	$[-1.81848, -1.00100] \times 10^{-3}$	14/14	first 14 amplitude levels
PLL and slow speed loop	$[-8.93731, -2.95865] \times 10^{-4}$	14/14	first 14 amplitude levels

Each of the two sets of 14 matrices, spanning $A_{\mathrm{exc}} = 0.0020496$ to 0.0314870, has negative spectral abscissa. The next level, $A_{\mathrm{exc}} = 0.0355189$, is not one of these linearized samples. In particular, the controlled endpoints near $A_{\mathrm{exc}} = 0.03641$ lie outside this finite-harmonic interval; successful terminal simulations are not covered by the spectral table. The matrix construction, the local unsaturated assumptions, and the projection-residual comparison are given in Section F.6.

7.6 Accompanying the response: result and scope

The two loops achieve different parts of resonant maintenance. The PLL adapts the frequency of the torque to regulate its phase lag relative to joint motion. The slower amplitude loop changes mean forward speed while leaving the within-cycle waveform to the mechanism. Together, they reach the same high-speed neighborhood from rest and from the named low-amplitude response, without consulting a stored family online. The ablation tests show why neither loop alone accomplishes both tasks.

This objective belongs to a wider family of mechanically informed control approaches. Orbital stabilization regulates deviations from a periodic motion rather than a complete time-indexed trajectory (Canudas-de Wit et al., 2002); hybrid zero dynamics exploits invariant structure in impact-driven walking (Westervelt et al., 2003). The continuously constrained model here has neither their virtual-constraint construction nor an impact/reset map. Likewise, task hierarchies and contact-force optimization solve a broader whole-body coordination problem (Mansard et al., 2009; Escande et al., 2014; Saab et al., 2013). Here the contribution is the local regulation of two response observables selected by the locomotor mechanics.

The simulations establish nominal phase–speed progression for the two specified initializations. A full-cycle comparison and final-stage step refinement support the endpoint interpretation; sampled finite-harmonic linearizations support the local controller design over the smaller declared amplitude range. They do not constitute a theorem of orbital stability for the exact nonlinear robot, global acquisition, or physical recovery after an unsuccessful transition. These are the next questions raised by the result, not properties inferred from phase lock alone.

The scientific link back to natural locomotion is the actuation philosophy: first identify motions organized by the mechanics, then supply energy with appropriate timing and adjust a locomotor

observable. The local same-mechanics bridge and the separately parameterized controlled response study demonstrate different parts of this programme; they are not a demonstrated controller navigating the entire conservative NLM.

Part IV

Synthesis, limits, and outlook

Chapter 8

Synthesis: accompanying the mechanics

8.1 What the question has become

The opening question asked which movements a robot's mechanics can favour, and how an actuator can sustain them. The answer is not one preferred waveform. It is a repertoire of coordinated motions, together with a principle that explains their energetic organization and methods for computing and maintaining them.

The opening cycle in Figure 1.2 is the simplest way to revisit that answer: an internal loop closes, while the robot's placement changes. The later family plots ask how many such rhythms the same mechanism can support; the maintained-response and control plots ask how timed energy input can make related motions usable under losses.

The ideal conservative model first separates organization from energetic maintenance. A natural locomotion cycle repeats its internal state while its body accumulates a displacement. This combines the coherent rhythm of a natural oscillation with the geometry of locomotion. It is not enough for a drawing of the joint motion to close, and it is not desirable for the entire world pose to close. Joint positions, joint velocities and the admissible body velocities must repeat together under the passive equations.

This distinction makes the nonlinear-mode analogy productive. A linear mode has a fixed frequency and choreography in its approximation near equilibrium. Nonlinear families permit the rhythm and coordination to change with energy. For a locomotor, environmental constraints add the coupling through which articulations and body motion participate in the same internal dynamics. The thesis combines these established ideas to identify an effective oscillatory channel, a propulsive channel and the energy transferred between them.

8.2 A principle with a precise scalar realization

The principle states that natural locomotion requires energy transferred between the effective oscillatory and propulsive channels to be returned over a cycle, without average gain or loss in either. This is the locomotor analogue of the repeated exchange between kinetic and potential energy in an ordinary oscillator. The partitions are different: the effective oscillator itself contains kinetic and potential energy. The side-by-side partitions in Figure 3.1 make that distinction visible: the analogy concerns cyclic redistribution, not a relabelling of the pendulum's kinetic and potential reservoirs.

In the declared conservative mechanics,

$$E = E_{\perp} + E_{\mathrm{prop}}, \qquad \dot{E_{\perp}} = P_{\mathrm{POE}} = -\dot{E_{\mathrm{prop}}}, \qquad I_{\mathrm{POE}}(T) = \int_0^T P_{\mathrm{POE}}(t)\ \mathrm{d}t. \tag{8.1}$$

The ideal constraints organize this redistribution without providing external power. The propulsive term is the complementary energy obtained from the completed-square decomposition, not the whole kinetic energy of the body and not a unique causal share of displacement.

The result is strongest in the one-joint 2SEG model. At a fixed joint angle crossed in a fixed direction, the effective potential and inertia are fixed. Recovering the effective oscillator's energy is therefore equivalent to recovering its joint speed. Once the body velocities also agree, this completes repetition of the internal state. Theorem 3.2 turns that physical observation into a necessary-and-sufficient characterization under explicit regular-section hypotheses. With nonzero translation, the cycles so characterized are natural locomotion cycles; regular neighbouring cycles and their phases form local NLM geometry.

For 3SEG, there are two joints. One energy value can be shared by different configurations and velocities at the observation section. The scalar POE balance remains necessary but cannot replace all the conditions for repetition. The method accordingly solves the full internal-state periodicity equations and then identifies in-phase and anti-phase coordination. This dimensional boundary is part of the contribution: it explains both the reach of the principle and the special strength of its scalar realization.

Neither result makes an existence claim for every mechanical design. The theorem characterizes cycles in its domain; the computation and design studies address where such cycles can actually be obtained. Likewise, conservation and periodicity do not alone imply attraction or stability under perturbations.

8.3 A gait repertoire and a mechanical-design question

A Natural Locomotion Manifold gives a robot designer more than one isolated gait. On a regular portion it has a cycle coordinate and a phase coordinate. Where mean forward body speed varies regularly from cycle to cycle, these can be written as $(\bar{v}, \varphi)$. The geometry then has a direct interpretation: choose an average advance, and locate the robot within the corresponding rhythm.

In Figure 4.2, moving around a plotted loop follows one cycle; comparing different loops changes the selected mean speed. The two operations are geometrically distinct. This is why a repertoire can guide a choice of gait without prescribing every point of its waveform.

The current two-segment evidence includes a connected numerical branch of 49 corrected cycles with independent periodicity, transport and local regularity checks at the sampled nodes. A separate set of 29 roots provides additional checks of the scalar construction. The three-segment study supplies 17 independently qualified primitive cycles, while its atlas uses a distinct population of 64 corrected supports to display larger-scale geometry. These collections serve different purposes; neither smooth interpolation nor a large historical export increases the number of independently verified motions.

The repertoire has two robotic uses. For an already specified mechanism, the problem is to calculate its available natural cycles. Before the mechanism is chosen, the problem is to adjust its mechanical parameters so that useful cycles can occur. The two-segment inertial relation $\gamma_\star(\varepsilon)$ provides an analytical landmark for the restricted design class studied. The corrected parameter-path examples approach it with small oscillations, but do not establish a universal existence boundary. The three-segment parameter studies similarly show that inertia, potential and geometry can materially affect the computed motions. They are structured design investigations, not a global optimizer or a proof that an unsuccessful solve means no cycle exists.

Moving between cycles is a further dynamical problem. In the passive ideal model, a state initialized on one periodic member remains on that member; it does not choose a new speed spontaneously. An actuator or another intervention is needed to change the selected motion. The conservative repertoire therefore informs control without already solving acquisition, transition or stabilization.

Transport also requires careful interpretation. Mean longitudinal body speed differs from world-axis displacement per period when the robot turns. Comparing an oscillatory gait with a matched

straight relative equilibrium can reveal changes in transport, but does not assign a coordinate-independent fraction of propulsion to oscillation. The energy-channel partition and the geometric displacement comparison answer different questions.

8.4 What the actuator should supply

In the real world, sustained positive losses require sustained energy input. For a periodic response of the damped and actuated model,

$$\int_0^T P_{\mathrm{in}}\ \mathrm{d}t = \int_0^T P_{\mathrm{diss}}\ \mathrm{d}t. \tag{8.2}$$

This is the total-energy work balance. It is not a new name for POE, which concerns transfer between internal channels. Both balances can be useful, but they should not be identified.

The local conservative-to-forced continuation makes energetic maintenance visible. Starting from one natural cycle, increasing dissipation and actuation yields a nearby synchronized response while preserving the robot's geometry, inertia and potential. Its period and waveform change: maintenance need not reproduce the conservative trajectory exactly. This numerical bridge illustrates a local deformation of one cycle, not persistence of an entire NLM under arbitrary losses.

The broader forced-response study uses different mechanical parameters and a different potential. It explores the synchronized locomotor responses associated with varying input amplitude and period. Regular local components of the transported solution set define the RLM in orbit–forcing space. This is not the same construction as a cycle–phase NLM in one conservative internal-state space. At fixed forcing phase, the reported RLM has two free forcing coordinates under the rank hypothesis; adding the forcing phase gives a three-dimensional lifted object. The common lesson is the organization of complete responses into families, not an identity of their dimensions or ambient spaces.

A swing provides the appropriate control image. The input should arrive at a useful time, replacing lost energy while allowing the mechanism to organize the motion. The controller studied here therefore adjusts two quantities instead of prescribing every joint angle. A phase-locked loop changes forcing frequency to regulate the joint-angle/torque fundamental phase lag. A slower amplitude loop regulates mean forward body speed.

At fundamental quadrature, the angle lags the torque by a quarter-period and the angular-velocity fundamental is in phase with the torque. This explains why the timing can supply positive mean work. It does not make quadrature a universal resonance or optimality criterion. Its association with a positive-peak-speed ridge is local to the reported lower-amplitude responses, and that peak-speed observable differs from the mean speed regulated by the controller.

This separates the jobs seen in Figure 7.3: the phase loop adjusts when energy is supplied, whereas the speed loop adjusts how strongly the actuator excites the motion. The mechanical response supplies the detailed waveform that neither loop prescribes.

The two nominal numerical traversals reach mean speeds near 12.068 for a target of 12, with small terminal phase error, from two distinct initializations. They demonstrate local phase–speed regulation without querying an offline joint-trajectory table. The sampled reduced-model stability calculations support local controller design, not a proof of exact nonlinear plant stability. Likewise, restoring a saved simulation state after a rejected test is not a recovery action available to a physical robot. Hardware acquisition and recovery remain separate tasks.

8.5 The next questions opened by the results

Three continuations of the work follow directly from these findings.

First, the local family geometry can be strengthened. The current two-segment branch has sampled checks of the periodic-orbit derivative and phase-swept geometry. Extending these checks along other branches, through relevant bifurcations, and across mechanical parameters would clarify where one speed coordinate remains useful and where a new chart is needed. In three segments, connecting independently corrected primitive witnesses with verified continuation would distinguish a collection of motions from a demonstrated connected branch. For the forced population, recovering full orbit-equation Jacobians is needed to test local RLM regularity systematically.

Second, the bridge between conservative organization and controlled maintenance can be developed across more than one cycle. Continuing both a family coordinate and dissipation would test when a useful portion of an NLM deforms into maintained responses, where folds or instabilities intervene, and how control can move between them. Floquet calculations, perturbation experiments in simulation and end-to-end recovery trials would address the dynamical properties that energy balance alone leaves open.

Third, identified physical mechanisms must test the modelling choices. Damping laws, actuator limits, sensing delays and environmental interaction can change the available cycles and the useful phase marker. Contact switching and impacts require hybrid or nonsmooth equations (Brogliato, 2016; Prasad et al., 2023). Extending the approach to fluid interaction would instead require appropriate forces or coupled fluid dynamics. The organizing question can carry over, but the present smooth ideal-constraint proof does not automatically carry over with it.

The thesis thus offers a principle, a scalar characterization, a computational and design route, and a local strategy for resonant maintenance. Its central proposal is to let the mechanics reveal an organized locomotor repertoire, then use actuation to sustain and regulate that motion. Control remains essential, but its role can be to accompany a rhythm that the robot helps determine.

Appendix A

A common reduction for continuously constrained locomotors

The two- and three-segment locomotors differ in dimension, but their equations are obtained by the same operation. This appendix presents that operation once, then identifies the model-specific ingredients used in appendices B and D. The output is an evolution equation for the admissible velocities, together with the kinematic equation that recovers configuration. Removing pose requires an additional symmetry, identified below; eliminating a reaction force does not by itself remove pose. The derivation applies to smooth, bilateral, ideal velocity constraints that remain continuously active and have constant rank. It does not cover impact resets, unilateral contact, stick–slip transitions, or changes of contact mode.

A.1 Admissible velocities and virtual displacements

Let $q \in Q$ be an n-dimensional configuration and consider k independent linear velocity constraints

$$A(q)\dot{q} = 0, \qquad \operatorname{rank} A(q) = k. \tag{A.1}$$

The admissible velocity space $\mathcal{D}_q = \ker A(q) \subset T_qQ$ therefore has dimension $m = n - k$. Choose a smooth full-column-rank matrix $H(q) \in \mathbb{R}^{n\times m}$ whose image is this kernel:

$$\dot{q} = H(q)\eta, \qquad \operatorname{Im} H(q) = \ker A(q), \qquad A(q)H(q) = 0. \tag{A.2}$$

The components of η are admissible quasi-velocities. They are not new configuration coordinates in general: for a nonintegrable distribution there need not exist a change of configuration variables whose time derivative is η.

In the Lagrange–d'Alembert model, the same distribution restricts virtual displacements to $\delta q = H(q)\delta\zeta$. Thus the matrix that parameterizes admissible velocities also supplies the directions on which the dynamical balance must hold. Here $\delta\zeta \in \mathbb{R}^m$ is an arbitrary virtual coefficient, not the variation of a newly introduced coordinate. Ideality means that the reaction does no work in any of these virtual directions.

A.2 Projected Lagrange–d'Alembert equation

For $K = \frac{1}{2}\dot{q}^{\mathsf{T}} M_q(q)\dot{q}$ with positive-definite inertia M_q, let $L = K - V$. Write $Q_{\mathrm{app}}(q, \dot{q}, t)$ for the applied generalized force (to distinguish it from the configuration space Q) and $\lambda \in \mathbb{R}^k$ for the constraint multiplier. The constrained balance is

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q_{\mathrm{app}} = A(q)^{\mathsf{T}}\lambda. \tag{A.3}$$

Multiplication by $H(q)^{\mathsf{T}}$ removes the unknown reaction because $H^{\mathsf{T}}A^{\mathsf{T}} = 0$:

$$H(q)^{\mathsf{T}}\left[\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q_{\mathrm{app}}\right] = 0. \tag{A.4}$$

After substituting $\dot{q} = H(q)\eta$, this equation has the reduced form

$$M_r(q)\dot{\eta} + h_r(q,\eta) + H(q)^{\mathsf{T}}\nabla V(q) = H(q)^{\mathsf{T}}Q_{\mathrm{app}}, \qquad M_r = H^{\mathsf{T}}M_qH. \tag{A.5}$$

To recover the inertial vector as well as the mass matrix, first form the unconstrained kinetic expression $M_q\ddot{q} + h_q(q,\dot{q})$, where $(h_q)_i = \sum_{j,k}(\partial_j(M_q)_{ik} - \frac{1}{2}\partial_i(M_q)_{jk})\dot{q}_j\dot{q}_k$. Then substitute $\ddot{q} = H\dot{\eta} + \dot{H}\eta$:

$$h_r = H^{\mathsf{T}}\{M_q\dot{H}\eta + h_q(q,H\eta)\}, \qquad \dot{H} = DH(q)[H(q)\eta]. \tag{A.6}$$

The two terms account for the moving admissible basis and the changing configuration-space inertia. Since H has independent columns, M_r is positive definite. Together, Equations (A.2) and (A.5) therefore give an ordinary differential equation on the admissible state space without solving for the reaction multiplier.

Figure A.1 summarizes the resulting sequence from the mechanism description to reduced integration and pose reconstruction. The first three boxes describe the mechanism and its admissible velocities; projection then determines their dynamics. The last box separates that calculation from placing the resulting motion in the plane, using the pose symmetry of the two models below.

A.3 Why ideal constraints redistribute but do not energize

Let $E = \dot{q}^{\mathsf{T}}\partial_{\dot{q}}L - L$ be the mechanical energy. For a time-independent Lagrangian, multiplication of Equation (A.3) by $\dot{q}^{\mathsf{T}}$ gives

$$\dot{E} = \dot{q}^{\mathsf{T}}Q_{\mathrm{app}} + \underbrace{\dot{q}^{\mathsf{T}}A(q)^{\mathsf{T}}\lambda}_{(A(q)\dot{q})^{\mathsf{T}}\lambda=0}. \tag{A.7}$$

The ideal constraint reaction therefore has zero total power. This does not make the constraint mechanically irrelevant. By selecting $\mathcal{D}_q$ and entering the reduced kinetic metric M_r, it structures how motion in one admissible direction couples to another. The effective oscillator–propulsive-channel split in Chapter 3 is a decomposition of that conservative coupling, not an external energy source.

When losses and actuation are introduced, only their applied powers appear:

$$Q_{\mathrm{app}} = Q_{\mathrm{loss}} + Q_{\mathrm{act}}, \qquad \dot{E} = \dot{q}^{\mathsf{T}}Q_{\mathrm{loss}} + \dot{q}^{\mathsf{T}}Q_{\mathrm{act}}. \tag{A.8}$$

This identity is the common mechanical origin of the cycle-integrated input–loss balance used in Chapter 6.

A.4 The two model instances

Table A.1 separates three dimensions that can otherwise be confused: configuration dimension, admissible-velocity dimension, and the number of transverse oscillatory directions exposed after carrier elimination. The last number is a property of the derived decomposition; it is not the total degree of freedom of the locomotor.

For the two-segment locomotor, the dimensional construction and normalization are given in appendix B. For the three-segment locomotor, Section D.1 instantiates the same projection with N in place of H. In both cases the Lagrangian and constraint are unchanged by a common planar rigid displacement of the mechanism. This symmetry makes the pose variables absent from the internal field after expressing translational velocities in the body frame. Once a reduced trajectory has been obtained, the group equation is integrated to recover the pose. That final integration is the pose reconstruction used throughout the manuscript.

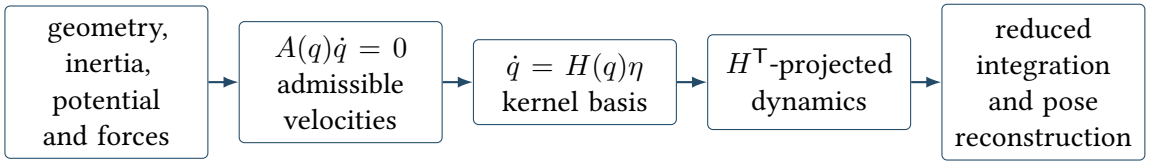


Figure A.1: Common modelling sequence for the two locomotors. The constraint matrix and kernel basis change with the mechanism, but projection eliminates the reaction multiplier in the same way. Pose reconstruction is the last operation; it is not another name for constraint reduction.

Table A.1: Configuration, admissible velocities and effective oscillatory directions count different objects. Both locomotors have two admissible body velocities; adding a joint adds one shape coordinate and one joint rate to the internal state.

	Two-segment locomotor	Three-segment locomotor
Configuration	$q = (x, y, \theta, \delta)$, $n = 4$	$q = (x, y, \theta, \delta_1, \delta_2)$, $n = 5$
Constraint rank	one no-side-slip row	one no-side-slip row
Admissible velocity	$\eta = (v, \omega, \sigma)$, $m = 3$	$\xi = (v, \omega, \sigma_1, \sigma_2)$, $m = 4$
Kernel basis	H in Equation (4.3)	N in Equation (D.2)
Transverse oscillatory directions	one after carrier elimination	two after carrier elimination
Return implication	a scalar section-storage value can identify the missing crossing speed	after a regular one-event section and carrier return, three endpoint coordinates remain; one storage value leaves generically a two-dimensional level surface, so full vector return remains necessary

A.5 Scope of the shared method

The projection in Equation (A.4) is general within its stated regularity class and is standard in nonholonomic mechanics (Bloch, 2015). Its use here does not establish a theory for all forms of environmental interaction. A walking model with impacts requires guards and reset maps; unilateral contact requires complementarity or contact mode logic; fluid locomotion requires explicit fluid forces or a coupled fluid–structure model. Those systems may preserve the organizing question of natural locomotion, but they do not inherit Equation (A.5) without new modelling work.

Appendix B

Two-segment model and conventions

This appendix gives the dimensional origin of the reduced equations used in Chapters 3 and 4. Its purpose is traceability: the main text works with the compact dimensionless system, while the equations below show which physical balance produces each coefficient and which sign convention is being used.

B.1 Geometry, velocities, and ideal constraint

Let m_1, I_1 be the mass and planar yaw inertia of the carrier and let m_2, I_2 denote the corresponding quantities for the internal segment. The distance from the contact point to the carrier centre of mass is l_1; the distance from the joint to the internal-segment centre of mass is l_2. The configuration and admissible velocity coordinates are

$$q = (x, y, \theta, \delta), \qquad \eta = (v, \omega, \sigma)^{\mathsf{T}}, \qquad \dot{q} = H(\theta)\eta, \tag{B.1}$$

with H given in Equation (4.3). Thus v is positive along the carrier heading, $\omega = \dot{\theta}$, and $\sigma = \dot{\delta}$. The absolute angular speed of the internal segment is $\omega + \sigma$. This last convention fixes the signs of all mixed inertia terms.

Let $r_C = (x, y)^{\mathsf{T}}$ be the constrained contact point and define

$$e(\varphi) = \begin{bmatrix} \cos\varphi \\ \sin\varphi \end{bmatrix}, \qquad e_\perp(\varphi) = \begin{bmatrix} -\sin\varphi \\ \cos\varphi \end{bmatrix}. \tag{B.2}$$

In the geometry used by the reduced model, the carrier centre of mass and the joint are co-located at

$$r_1 = r_J = r_C + l_1 e(\theta), \qquad r_2 = r_J + l_2 e(\theta + \delta), \tag{B.3}$$

where r_2 is the internal-segment centre of mass. Since $\dot{r}_C = v e(\theta)$,

$$\dot{r}_1 = v e(\theta) + l_1 \omega e_\perp(\theta), \tag{B.4}$$

$$\dot{r}_2 = \dot{r}_1 + l_2(\omega + \sigma) e_\perp(\theta + \delta). \tag{B.5}$$

These positions remove any ambiguity about the reference points of l_1 and l_2. They also make the constraint physical: the contact point may move along $e(\theta)$ but not along $e_\perp(\theta)$. Its reaction need not vanish, although its power on this admissible motion is zero.

The constrained Lagrange–d'Alembert equation is

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = A(\theta)^{\mathsf{T}}\lambda, \qquad L = K - U(\delta). \tag{B.6}$$

Because $A(\theta)H(\theta) = 0$, multiplication by H^{T} eliminates the reaction multiplier. The ideal lateral reaction consequently performs no work on an admissible velocity. The projected balance is

$$M_{\mathrm{dim}}(\delta)\dot{\eta} + h_{\mathrm{dim}}(\delta, \eta) + \begin{bmatrix} 0 & 0 & U'(\delta) \end{bmatrix}^{\mathsf{T}} = 0. \tag{B.7}$$

B.2 Dimensional reduced equations

Set $M_{\mathrm{tot}} = m_1 + m_2$ and $D = \omega + \sigma$. The kinetic energy is

$$K = \frac{1}{2} m_1 \left\| \dot{r}_1 \right\|^2 + \frac{1}{2} I_1 \omega^2 + \frac{1}{2} m_2 \left\| \dot{r}_2 \right\|^2 + \frac{1}{2} I_2 D^2. \tag{B.8}$$

Direct evaluation and projection give

$$M_{\mathrm{dim}}(\delta) = \begin{bmatrix} M_{\mathrm{tot}} & -m_2 l_2 \sin\delta & -m_2 l_2 \sin\delta \\ -m_2 l_2 \sin\delta & I_1 + I_2 + M_{\mathrm{tot}} l_1^2 + m_2 l_2^2 + 2 m_2 l_1 l_2 \cos\delta & I_2 + m_2 l_2^2 + m_2 l_1 l_2 \cos\delta \\ -m_2 l_2 \sin\delta & I_2 + m_2 l_2^2 + m_2 l_1 l_2 \cos\delta & I_2 + m_2 l_2^2 \end{bmatrix}. \tag{B.9}$$

The quadratic inertial vector is

$$h_{\mathrm{dim}} = \begin{bmatrix} -M_{\mathrm{tot}} l_1 \omega^2 - m_2 l_2 \cos\delta \, D^2 \\ (M_{\mathrm{tot}} l_1 + m_2 l_2 \cos\delta) v\omega - m_2 l_1 l_2 \sin\delta \, (D^2 - \omega^2) \\ m_2 l_2 \cos\delta \, v\omega + m_2 l_1 l_2 \sin\delta \, \omega^2 \end{bmatrix}. \tag{B.10}$$

The first row is the longitudinal balance, the second is the carrier-yaw balance, and the third is the internal-joint balance. In particular, the finite-speed product $v\omega$ is present in the mechanical model before the change of variables $q_y = v\omega$ used in Appendix C.

The signs in Equations (B.9) and (B.10) can be checked without redoing the full variational calculation. At $\delta = 0$ the translation–rotation inertia terms proportional to $\sin\delta$ vanish. At $\omega = \sigma = 0$ the entire drift vector vanishes. Hence every state $(0, v, 0, 0)$ is a reduced equilibrium, exactly as required by Equation (4.16).

B.3 Dimensionless dictionary and energy check

Let

$$L = l_1 + l_2, \qquad \alpha = \frac{l_1}{L}, \qquad \beta = \frac{l_2}{L}, \qquad \alpha + \beta = 1, \tag{B.11}$$

and let t_* be the reference time. Hats denote dimensionless variables:

$$\widehat{x} = \frac{x}{L}, \quad \widehat{t} = \frac{t}{t_*}, \quad \widehat{v} = \frac{t_* v}{L}, \quad \widehat{\omega} = t_* \omega, \quad \widehat{\sigma} = t_* \sigma, \quad \widehat{E} = \frac{t_*^2 E}{M_{\mathrm{tot}} L^2}. \tag{B.12}$$

Because the joint angle δ is already dimensionless, dimensional spring coefficients in $U_{\mathrm{dim}} = \frac{1}{2} k_{2,\mathrm{dim}} \delta^2 + \frac{1}{4} k_{4,\mathrm{dim}} \delta^4$ become

$$\widehat{k}_2 = \frac{t_*^2 k_{2,\mathrm{dim}}}{M_{\mathrm{tot}} L^2}, \qquad \widehat{k}_4 = \frac{t_*^2 k_{4,\mathrm{dim}}}{M_{\mathrm{tot}} L^2}. \tag{B.13}$$

The numerical chapters drop hats after this normalization. Their periods, velocities, and energies are therefore dimensionless unless an explicit dimensional calibration is supplied; in particular, the reported period range and the values $k_2 = 1$, $k_4 = 10$ belong to this numerical scale. When $k_{2,\mathrm{dim}} > 0$, choosing $t_* = \sqrt{M_{\mathrm{tot}} L^2 / k_{2,\mathrm{dim}}}$ sets $\widehat{k}_2 = 1$; the computations may equivalently be read as fixing that reference clock.

The remaining design ratios are

$$\mu = \frac{m_2}{M_{\mathrm{tot}}}, \qquad j_i = \frac{I_i}{M_{\mathrm{tot}} L^2}, \tag{B.14}$$

and the design used in the conservative computations specializes to

$$\alpha = \varepsilon, \qquad \beta = 1 - \varepsilon, \mu = 1 - \varepsilon, \qquad j_1 = \gamma \alpha^3, \qquad j_2 = \gamma \beta^3. \tag{B.15}$$

Dividing Equation (B.9) by the common mass–length scales then produces

$$\frac{m_2 l_2}{M_{\mathrm{tot}} L} = \mu\beta, \qquad \frac{I_1}{M_{\mathrm{tot}} L^2} = j_1, \qquad \frac{I_2}{M_{\mathrm{tot}} L^2} = j_2, \tag{B.16}$$

and hence the functions $\rho, B_{11}, B_{12}, B_{22}$ in Equation (4.12). The common minus sign in the first row and column of Equation (4.11) is inherited from the chosen positive direction of δ; it is not absorbed into the definition of ρ.

For the unforced conservative model, the projected energy is

$$E = \frac{1}{2}\eta^{\mathsf{T}} M(\delta)\eta + U(\delta). \tag{B.17}$$

The dimensional and normalized inertial vectors obey $\eta^{\mathsf{T}} h = \frac{1}{2}\eta^{\mathsf{T}}\dot{M}\eta$, with $\dot{M} = M'(\delta)\sigma$. This is the power identity inherited from the ideal projection, even though the quasi-velocities are not coordinate derivatives. Differentiating and substituting Equation (4.10) gives

$$\dot{E} = \eta^{\mathsf{T}}\left(M\dot{\eta} + h + \begin{bmatrix}0 & 0 & U'\end{bmatrix}^{\mathsf{T}}\right) = 0, \tag{B.18}$$

The identity is analytical; numerical energy checks in Appendix C evaluate this conservation law independently of the nonlinear solver. It concerns total mechanical energy. The Schur-complement storage $E_\perp$ in Chapter 3 is a different scalar: its derivative need not vanish after the propulsive channel is reopened.

B.4 Linearization for the inertial design relation

At the straight state $(\delta, v, \omega, \sigma) = (0, s_0, 0, 0)$, the longitudinal equation has no term linear in a perturbation. The transverse coordinates $(\delta_1, \omega_1, \sigma_1)$ therefore satisfy

$$\dot{\delta}_1 = \sigma_1, \qquad B\begin{bmatrix}\dot{\omega}_1 \\ \dot{\sigma}_1\end{bmatrix} + s_0\omega_1\begin{bmatrix}\mathfrak{c} \\ \mathfrak{d}\end{bmatrix} + \begin{bmatrix}0 \\ k_2\delta_1\end{bmatrix} = 0. \tag{B.19}$$

Here B, $\mathfrak{c}$ and $\mathfrak{d}$ are the aligned coefficients of Equation (4.38). Inverting B gives

$$J_\perp(s_0) = \begin{bmatrix}0 & 0 & 1 \\ k_2 B_{12}/D_B & -s_0 C_2/D_B & 0 \\ -k_2 B_{11}/D_B & -s_0 N/D_B & 0\end{bmatrix}. \tag{B.20}$$

At $N = 0$, the joint pair (δ_1, σ_1) has characteristic polynomial $\lambda^2 + k_2 B_{11}/D_B$; yaw has the remaining eigenvalue $-s_0 C_2/D_B$. Equivalently,

$$\det(\lambda I - J_\perp) = \left(\lambda + \frac{s_0 C_2}{D_B}\right)\left(\lambda^2 + \frac{k_2 B_{11}}{D_B}\right) \quad \text{when } N = 0. \tag{B.21}$$

This proves the factorization used in the body. The cubic spring coefficient k_4 does not enter this linearization, but remains in every nonlinear periodic-orbit computation.

With $\beta = 1 - \varepsilon$, substitution of Equation (4.40) into $N = B_{11}\beta^2 - B_{12}(\varepsilon + \beta^2)$ gives

$$N = \varepsilon\beta^2\{\gamma(\varepsilon^2 + \varepsilon - 1) + \varepsilon(\varepsilon - 1)\}.$$

Its vanishing yields $\gamma_\star$ in Equation (4.42). The factorization predicts the clock at that relation. It does not establish the existence or continuation of nonlinear cycles on either side of it; the fourteen corrected design samples provide the separate nonlinear evidence.

B.5 Conventions for the numerical results

The two fixed-design datasets in the body have different roles. The 29 repolished scalar roots check the POE formulation on both nonzero-speed sheets. The 49-node predecessor-connected continuation examines one positive-speed family and its local cycle–phase geometry. Each uses the same reference mechanics Equation (4.9); the design study instead changes γ. The independent integration and rank checks are specified in Appendix C.

The dimensional model and scaling above apply to all these computations. A normalized speed is not a calibrated hardware speed. In particular, $\bar{v}$ averages the longitudinal body velocity, whereas $\Delta x/T$ also depends on heading over the cycle. Exact internal periodicity concerns $(\delta, v, \omega, \sigma)$; pose is reconstructed separately. When damping and torque are introduced later, their work appears explicitly in the energy balance rather than being absorbed into POE.

Appendix C

Scalar coordinates and reproducible periodicity checks

This appendix connects the physical four-state model to the coordinates used to compute its cycles, then specifies the checks on the 29 scalar roots and the distinct 49-node connected family of Chapter 4. The changes of coordinates neither prescribe a joint waveform nor approximate the mechanics. In particular, the derivative of effective inertia belongs in the storage-rate identity; adding its auxiliary-oscillator term to the original shape acceleration would change the model being solved.

C.1 Coefficient functions and opened field

Start from the physical internal state $z = (\delta, v, \omega, \sigma)$ of Appendix B. The coefficient $r(\delta)$ below is the function denoted r_q in Equation (3.22); it is not a shape coordinate. The inertia functions B_{ij}, ρ, Δ and $M_\perp$ are those of Equation (4.12) and Section 3.2.

For the dimensionless parameters introduced in Chapters 3 and 4, define

$$c_{\mathrm{geom}}(\delta) = \alpha + \mu\beta\cos\delta, \tag{C.1}$$

$$A_q(\delta) = \alpha j_2 + \alpha\beta^2\mu(1-\mu) - \mu\beta j_1\cos\delta, \tag{C.2}$$

$$r(\delta) = \frac{A_q(\delta)}{\Delta(\delta)}, \tag{C.3}$$

$$E_{\mathrm{geom}}(\delta) = B_{11}(\delta)\big(c_{\mathrm{geom}}(\delta) - \alpha\big) + \alpha\rho(\delta)^2. \tag{C.4}$$

With the scaled state

$$y = (\delta, \sigma, \nu, q_y), \qquad h = |\bar{v}|^{-1}, \qquad \nu = hv, \qquad q_y = v\omega, \tag{C.5}$$

the physical velocities are

$$v = \frac{\nu}{h}, \qquad \omega = \frac{hq_y}{\nu}. \tag{C.6}$$

Here h is a fixed label during each solve, ν scales longitudinal velocity by its requested mean magnitude, and q_y collects the product $v\omega$ already present in the physical balances. The chart requires $h > 0$ and $\nu \neq 0$: a nonzero mean alone would not exclude an instantaneous zero of v. The shape acceleration is

$$F_\sigma = \frac{r\{q_y - \rho(\sigma+\omega)^2\} - U'}{M_\perp}. \tag{C.7}$$

The two transport rows are evaluated in a fixed dependency order. First set

$$\lambda_{\mathrm{tr}} = \frac{c_{\mathrm{geom}}}{\Delta}\left(\nu - 2h\rho\sigma\right), \tag{C.8}$$

$$\phi_{\mathrm{tr}} = \nu\left[-\frac{B_{12}-\rho^2}{\Delta}F_\sigma + \frac{\rho c_{\mathrm{geom}}}{\Delta}\sigma^2\right], \tag{C.9}$$

$$R_{\mathrm{tr}} = \frac{\rho(B_{11} - B_{12})F_\sigma + E_{\mathrm{geom}}(2\omega\sigma + \sigma^2) + B_{11}c_{\mathrm{geom}}\omega^2}{\Delta}. \tag{C.10}$$

Then

$$F_q = \frac{\phi_{\mathrm{tr}} + h^2\nu^{-1}q_y R_{\mathrm{tr}} - \lambda_{\mathrm{tr}}q_y}{h}, \tag{C.11}$$

and

$$F_\nu = \frac{h\rho F_\sigma + h(c_{\mathrm{geom}} - \alpha)\sigma^2 + h^2\rho\nu^{-1}F_q + 2h^2(c_{\mathrm{geom}} - \alpha)\nu^{-1}q_y\sigma + h^3c_{\mathrm{geom}}\nu^{-2}q_y^2}{1 + h^2\rho\nu^{-2}q_y}. \tag{C.12}$$

The denominator

$$D_{\mathrm{fs}} = 1 + h^2\rho\nu^{-2}q_y \tag{C.13}$$

is monitored along every returned orbit.

The exact scaled system is

$$\dot{y} = (\sigma, F_\sigma, F_\nu, F_q). \tag{C.14}$$

For $s = \mathrm{sgn}(\bar{v})$, the oriented variables

$$Y = (\delta, \zeta, u, Q), \qquad (\sigma, \nu, q_y) = (s\zeta, su, sQ) \tag{C.15}$$

and time change $\mathrm{d}/\,\mathrm{d}\vartheta = s\ \mathrm{d}/\,\mathrm{d}t$ give

$$\frac{\mathrm{d}Y}{\mathrm{d}\vartheta} = (\zeta, F_\sigma, F_\nu, F_q), \tag{C.16}$$

where the three field components on the right are evaluated at the signed scaled state. The inverse is $(v, \omega, \sigma) = (su/h, hQ/u, s\zeta)$. Thus $u > 0$ describes one sign-definite speed sheet and both signs use the same positive joint crossing in oriented time. For $s = -1$, that clock runs backward relative to physical time. This is an exact coordinate and time transformation, not a large-speed approximation.

C.2 Exchange identity in the implemented convention

The conservative transverse field is

$$F_{\mathrm{cl}} = -\frac{U' + \frac{1}{2}M_\perp'\sigma^2}{M_\perp}, \tag{C.17}$$

whereas the opened field is Equation (C.7). Their difference is

$$M_\perp(F_\sigma - F_{\mathrm{cl}}) = r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M_\perp'\sigma^2. \tag{C.18}$$

Consequently,

$$P_{\mathrm{POE}} = \sigma M_\perp(F_\sigma - F_{\mathrm{cl}}) = \sigma\left[r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M_\perp'\sigma^2\right]. \tag{C.19}$$

Direct differentiation of

$$E_\perp = \frac{1}{2}M_\perp(\delta)\sigma^2 + U(\delta) \tag{C.20}$$

along the opened field gives

$$\frac{\mathrm{d}E_\perp}{\mathrm{d}t} = \sigma\left(M_\perp F_\sigma + U' + \frac{1}{2}M_\perp'\sigma^2\right) = P_{\mathrm{POE}}. \tag{C.21}$$

Under the oriented time transformation, the accumulator uses

$$\frac{\mathrm{d}I_{\mathrm{POE}}}{\mathrm{d}\vartheta} = sP_{\mathrm{POE}}, \tag{C.22}$$

so it integrates exchange along the physical-time traversal represented by the oriented orbit. For $s = +1$ this is the forward-period integral. For $s = -1$, increasing ϑ traverses physical time backward, and the result is the backward-period integral. Reversing a complete periodic traversal changes its sign but not its zero set; this is the only orientation-invariant fact used by the return theorem.

The scalar storage identifies a missing joint rate only after choosing a common angle and crossing direction. On $\delta = 0, \zeta > 0, E_\perp = \frac{1}{2}M_\perp(0)\zeta^2 + U(0)$ is strictly increasing in ζ when $M_\perp(0) > 0$. Hence zero integrated POE fixes the terminal ζ there. The additional conditions on u and Q fix the body velocities through the inverse map above. This is the conditional energy–speed identification used by the physical theorem, not conservation of $E_\perp$ along every passive trajectory.

C.3 The scalar solve and its independent checks

At requested mean velocity $\bar{v} \neq 0$, logarithmic unknowns

$$p = (a, b, Q_0), \qquad Y_0 = (0, e^a, e^b, Q_0) \tag{C.23}$$

preserve the positive oriented joint crossing and the positive scaled speed. An initial dead interval excludes the starting event itself. The subsequent same-direction crossing defines the duration. The equations also integrate the two accumulators

$$\frac{\mathrm{d}I_{\mathrm{POE}}}{\mathrm{d}\vartheta} = sP_{\mathrm{POE}}, \qquad \frac{\mathrm{d}I_u}{\mathrm{d}\vartheta} = u. \tag{C.24}$$

The four residual components in Equation (4.26) use

$$e_{\mathrm{sc}} = 1, \qquad u_{\mathrm{sc}} = \max(1, |u_0|), \qquad Q_{\mathrm{sc}} = \max(1, |Q_0|). \tag{C.25}$$

The POE row has a fixed unit scale in these normalized coordinates, not a scale recomputed from its current error. All components are retained by the least-squares solve. The first three have one dependence imposed by the conserved total energy; the speed condition supplies the third independent direction.

The reported realization re-polishes 29 stored section states. The stored states are predictors, not accepted endpoints. A local least-squares solve uses DOP853 with relative and absolute tolerances 10^{-10} and 10^{-12}, maximum step $T_{\mathrm{ref}}/300$, and at most 12 residual evaluations. Its result is independently integrated at nominal tolerances $(10^{-11}, 10^{-13})$ and refined tolerances $(2 \times 10^{-12}, 2 \times 10^{-14})$, with maximum steps $T_{\mathrm{ref}}/500$ and $T_{\mathrm{ref}}/1000$. The nominal and refined domain diagnostics sample 2,001 and 4,001 states. The initial event guard is $0.05T_{\mathrm{ref}}$ and the search horizon is $2T_{\mathrm{ref}}$. Here T_{ref} is the predictor's period scale, not a period prescribed to the final solution.

Acceptance checks the event and the complete terminal state, the mean-speed condition, total-energy conservation, the POE identity and independently reconstructed pose. Positivity and nonsingularity are evaluated along the trajectory: $\Delta > 0$, $M_\perp > 0$, u away from zero and the finite-speed denominator away from zero. The numerical protocol uses the exact original field, rather than a periodically closed spline as a substitute for its solution.

Pose integration uses the physical forward-time convention even when the oriented chart traverses time backward. Refinement changes the reconstructed pose by at most 4.26×10^{-14} in the reported comparison. A nonidentity increment and a resolved matched-straight deviation are observed for every retained root. Refinement agreement is a numerical convergence diagnostic, not an independent proof of global accuracy.

C.3.1 Rank and conservative dependence

Differentiate in $(\log \zeta_0, \log u_0, Q_0)$, with unit scales on the logarithmic columns and scale $\max(1, |Q_0|)$ on the Q_0 column. When $\bar{v}$ is also varied, its scale is $\max(1, |\bar{v}|)$. Central differences with relative steps 3×10^{-6}, 10^{-5} and 3×10^{-5} evaluate the return-map Jacobians, using nominal integration tolerances $(10^{-11}, 10^{-13})$. A relative singular-value threshold of 10^{-7} defines the reported numerical rank. Residual row scales are frozen

Table C.1: Independent scalar-root checks on the 29 re-polished solutions. The state-defect norm refers to the normalized oriented coordinates Y. These finite numerical values are not rigorous interval error bounds.

Quantity	Observed value or interval
Accepted scalar roots	29
Maximum re-polished scaled residual	4.57×10^{-15}
Maximum independent terminal-state defect	2.85×10^{-14}
Maximum relative total-energy span	5.03×10^{-15}
Minimum Δ	1.4367497
Minimum $M_\perp$	0.09143638
Sampled u	$[0.5624729,\ 1.4384818]$
Finite-speed denominator	$[0.9638301,\ 1.1016363]$
Maximum storage–POE identity difference	1.07×10^{-14}
Maximum propulsive-storage–POE identity difference	1.09×10^{-12}

at the sampled root during these differences, so changing the perturbation does not change the norm in which the Jacobian is measured.

At all 29 solutions, the return-only 3×3 block has rank two. Adding the fixed-speed condition produces rank three, with

$$\sigma_{\min}(D_p \mathcal{R}_{\mathrm{POE}}) \in [0.021917,\ 0.104764], \qquad \kappa_2 \le 48.7703. \tag{C.26}$$

Allowing speed to vary gives a rank-three 4×4 matrix, whose smallest nonzero singular value lies in $[0.025463,\ 0.107124]$. At the nominal difference step the largest null-to-nonzero singular-value ratio is 4.46×10^{-9}; the rank pattern persists across all three steps. The exact energy relation explains the dependence and these numerical values resolve the remaining directions. They support local branches at the sampled roots, not a connection through zero speed.

C.4 A predecessor-connected 49-node family

The first positive-speed root initializes a distinct computation that solves the three independent conditions of Equation (4.28). Here the fixed storage and Q scales are 1 and 8, respectively; the unknowns are $(\log \zeta_0, \log u_0, Q_0/8)$. These scales differ from the root-dependent Q scale used for the 29-root calculation above. The next requested speed uses the preceding tangent predictor and Newton correction, with at most seven iterations and five backtracking reductions. The residual and step tolerances are 2×10^{-12} and 2×10^{-11}. The tangent is obtained from event-corrected variational equations. Independent finite differences of the complete event map at five distributed nodes check that derivative.

For a section defined by $h_\Sigma(z) = 0$, a fixed-time perturbation $\delta z(T)$ must be corrected for its event-time change:

$$\delta T = -\frac{Dh_\Sigma(z(T))\,\delta z(T)}{Dh_\Sigma(z(T))\,f(z(T))}, \qquad \delta z_{\mathrm{event}} = \delta z(T) + f(z(T))\,\delta T. \tag{C.27}$$

The denominator is nonzero at a transverse crossing. This standard event-sensitivity relation distinguishes differentiation of a flow at fixed time from differentiation of its return map.

At each node, separate nominal and refined integrations use tolerances $(10^{-10}, 10^{-12})$ and $(2 \times 10^{-11}, 2 \times 10^{-13})$, maximum steps $T_{\mathrm{ref}}/300$ and $T_{\mathrm{ref}}/600$, and a $0.04T_{\mathrm{ref}}$ initial guard. A further first-return calculation checks the primitive same-direction event. Pose is integrated from identity. The phase sweep samples 64 phases per cycle, and its flow direction is compared with the branch tangent.

The energy derivative and noncancellation check justify recovering the omitted u condition locally, while independent integration checks it directly as well. The nonzero speed tangent supports $\bar{v}$ as a local coordinate. Independence from the phase direction supports the sampled two-dimensional sweep. A predecessor-connected computation and a finite regularity check provide stronger evidence than an interpolated display alone, but do not certify all unsampled points or global embedding.

Table C.2: Checks on the connected positive-speed family. Every node is corrected from its predecessor under unchanged mechanical parameters.

Quantity	Observed value
Corrected, predecessor-connected nodes	49
Mean-speed interval	$[1.132791,\ 25]$
Maximum shooting residual norm	3.56×10^{-13}
Maximum refined scaled state defect	6.53×10^{-13}
Minimum shooting-Jacobian singular value	0.024776
Maximum shooting-Jacobian condition number	40.648
Minimum $\lvert\partial E/\partial u\rvert$	1.2864557
Minimum normalized noncancellation margin	0.3128472
Minimum normalized speed-tangent component	0.9929044
Minimum sampled branch–flow singular value	0.7610252
Maximum independent derivative discrepancy	1.79×10^{-9}

C.5 The inertia study uses the same physical tests

The design study reuses historical parameter-path points only as predictors. Fourteen selected points, seven on each of the two opposite-speed paths, are corrected with the exact passive field and independently integrated. Their mechanics changes through γ, so they are not extra members of the fixed-design 49-node NLM. Every reported point satisfies its cycle-level checks. None lies at or below $\gamma_\star$.

The near-limit amplitude fit uses four preselected points per path. For the positive-speed path, the through-origin coefficient in $\gamma - \gamma_\star = CA^2$ is 0.60344167, with $R^2 = 0.99999685$; free-intercept and leave-one-out comparisons support that local observation. The negative-speed fit is more sensitive to the point selection. Consequently the common quantitative-fit condition across both paths is not satisfied, even though the fourteen corrected cycles are retained. The body reports the positive-path fit and the qualitative amplitude/period trends separately.

All checks in this appendix concern normalized simulations of the declared conservative model. Their purpose is to connect a mathematical criterion to reproducible passive trajectories, not to claim hardware validation, unseeded discovery, global uniqueness, stability or computational superiority over other periodic-orbit methods.

Appendix D

Three-segment mechanics and periodic-cycle protocol

This appendix specifies the model and numerical procedures behind Chapter 5. The seventeen-sample local-family study and the sixty-four-support visual atlas are treated separately. The exact mechanical definitions are common; the numerical properties established for each set are not interchangeable.

D.1 Exact mechanical model

Let

$$e(\varphi) = \begin{bmatrix} \cos\varphi \\ \sin\varphi \end{bmatrix}, \qquad e^{\perp}(\varphi) = \begin{bmatrix} -\sin\varphi \\ \cos\varphi \end{bmatrix}. \tag{D.1}$$

The configuration is $q = (x, y, \theta, \delta_1, \delta_2)$, where $r_C = (x, y)^{\mathsf{T}}$ is the no-side-slip contact point. The constraint and an admissible velocity basis are

$$\begin{bmatrix} -\sin\theta & \cos\theta & 0 & 0 & 0 \end{bmatrix} \dot{q} = 0, \qquad \dot{q} = N(q)\xi, \qquad N(q) = \begin{bmatrix} \cos\theta & 0 & 0 & 0 \\ \sin\theta & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{D.2}$$

with $\xi = (v, \omega, \sigma_1, \sigma_2)^{\mathsf{T}}$. Thus $\dot{x} = v\cos\theta$, $\dot{y} = v\sin\theta$, $\dot{\theta} = \omega$, and $\dot{\delta}_i = \sigma_i$. Write $\boldsymbol{\delta} = (\delta_1, \delta_2)^{\mathsf{T}}$. The independent mechanical state is $(\boldsymbol{\delta}, \xi)$; N supplies pose reconstruction after that state has been integrated.

Set $\ell_2 = L_2/2$ and $\ell_3 = L_3/2$. The three body centers are

$$\begin{aligned} r_1 &= r_C + l_1 e(\theta), \\ r_2 &= r_1 + \ell_2 e(\theta + \delta_1), \\ r_3 &= r_1 + L_2 e(\theta + \delta_1) + \ell_3 e(\theta + \delta_1 + \delta_2), \end{aligned} \tag{D.3}$$

and their velocities are

$$\begin{aligned} \dot{r}_1 &= v e(\theta) + l_1 \omega e^{\perp}(\theta), \\ \dot{r}_2 &= \dot{r}_1 + \ell_2(\omega + \sigma_1) e^{\perp}(\theta + \delta_1), \\ \dot{r}_3 &= \dot{r}_1 + L_2(\omega + \sigma_1) e^{\perp}(\theta + \delta_1) \\ &\quad + \ell_3(\omega + \sigma_1 + \sigma_2) e^{\perp}(\theta + \delta_1 + \delta_2). \end{aligned} \tag{D.4}$$

The corresponding angular rates are $\Omega_1 = \omega$, $\Omega_2 = \omega + \sigma_1$, and $\Omega_3 = \omega + \sigma_1 + \sigma_2$. The kinetic energy used in the calculation is exactly

$$\mathcal{T} = \frac{1}{2}\sum_{i=1}^{3} m_i \|\dot{r}_i\|^2 + \frac{1}{2}\sum_{i=1}^{3} I_i \Omega_i^2 = \frac{1}{2}\xi^{\mathsf{T}} M(\boldsymbol{\delta})\xi. \tag{D.5}$$

This also specifies every entry of M without a long expanded matrix: if $\dot{r}_i = J_i\xi$ and $\Omega_i = b_i\xi$ are read from Equation (D.4), then $M = \sum_i (m_i J_i^{\mathsf{T}} J_i + I_i b_i^{\mathsf{T}} b_i)$. Here J_i is 2×4 and b_i is 1×4. Together with

$$\begin{aligned} U(\boldsymbol{\delta}) = &\frac{1}{2}k_{2,1}\delta_1^2 + \frac{1}{4}k_{4,1}\delta_1^4 + \frac{1}{2}k_{2,2}\delta_2^2 + \frac{1}{4}k_{4,2}\delta_2^4 \\ &+ \frac{1}{2}k_{12}(\delta_1 - \delta_2)^2, \end{aligned} \tag{D.6}$$

this defines the Lagrangian $L = \mathcal{T} - U$ without an implicit geometric convention. The projected Lagrange–d'Alembert equations are

$$N(q)^{\mathsf{T}}\left[\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q\right] = 0, \qquad Q = (0, 0, 0, u_1, u_2)^{\mathsf{T}}. \tag{D.7}$$

Equivalently,

$$M(\boldsymbol{\delta})\dot{\xi} + h(\boldsymbol{\delta}, \xi) + \begin{bmatrix} 0 \\ 0 \\ \partial_{\delta_1} U \\ \partial_{\delta_2} U \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ u_1 \\ u_2 \end{bmatrix}. \tag{D.8}$$

For completeness, the velocity-dependent vector in Equation (D.8) can be written compactly by defining

$$m_\Sigma = m_1 + m_2 + m_3, \quad a = \ell_2 m_2 + L_2 m_3, \quad D_1 = \omega + \sigma_1, \quad D_2 = \omega + \sigma_1 + \sigma_2, \tag{D.9}$$

and $c_i = \cos\delta_i$, $s_i = \sin\delta_i$, $c_{12} = \cos(\delta_1 + \delta_2)$, and $s_{12} = \sin(\delta_1 + \delta_2)$. Its four components are

$$\begin{aligned} h_v = &-l_1 m_\Sigma \omega^2 - a c_1 D_1^2 - \ell_3 m_3 c_{12} D_2^2, \\ h_\omega = &\,\omega v(l_1 m_\Sigma + a c_1 + \ell_3 m_3 c_{12}) - l_1 a(2\omega\sigma_1 + \sigma_1^2)s_1 \\ &- L_2 \ell_3 m_3 (2\omega\sigma_2 + 2\sigma_1\sigma_2 + \sigma_2^2)s_2 \\ &- l_1 \ell_3 m_3 \left[2\omega(\sigma_1 + \sigma_2) + (\sigma_1 + \sigma_2)^2\right] s_{12}, \\ h_{\sigma_1} = &\,\omega v(a c_1 + \ell_3 m_3 c_{12}) + \omega^2(l_1 a s_1 + l_1 \ell_3 m_3 s_{12}) \\ &- L_2 \ell_3 m_3 (2\omega\sigma_2 + 2\sigma_1\sigma_2 + \sigma_2^2)s_2, \\ h_{\sigma_2} = &\,\ell_3 m_3 \left(L_2 D_1^2 s_2 + l_1 \omega^2 s_{12} + \omega v c_{12}\right). \end{aligned} \tag{D.10}$$

In the state order used by the cycle files,

$$\dot{z} = \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ M(\boldsymbol{\delta})^{-1}\left(\begin{bmatrix} 0 \\ 0 \\ u_1 \\ u_2 \end{bmatrix} - h(\boldsymbol{\delta}, \xi) - \begin{bmatrix} 0 \\ 0 \\ \partial_{\delta_1} U \\ \partial_{\delta_2} U \end{bmatrix}\right) \end{bmatrix}. \tag{D.11}$$

For the conservative AP/IP calculations, $u_1 = u_2 = 0$, and the energy in Equations (D.5) and (D.6) is conserved by the exact field. The closed-support calculation evaluates the Schur complement and its shape derivatives from this same matrix M, then integrates the Euler–Lagrange field Equations (5.9) and (5.10); it is not inferred from energy conservation alone.

The same construction yields the instantaneous exchange, not only its endpoint balance. Denote the two joint-acceleration components of the original field by F_{op} and those of the auxiliary field by F_{cl}. Their difference defines the effective shape force and power

$$Q_{\mathrm{POE}} = M_\perp(F_{\mathrm{op}} - F_{\mathrm{cl}}), \qquad P_{\mathrm{POE}} = w^{\mathsf{T}} Q_{\mathrm{POE}} = \frac{\mathrm{d}E_\perp}{\mathrm{d}t}. \tag{D.12}$$

The last equality follows by differentiating the quadratic storage and using the auxiliary Euler–Lagrange equation. Positive power means that the effective oscillator receives energy; in the unforced model the propulsive channel loses the same amount.

D.2 Modal references and conservative continuation

The shared parameter vector is given, alongside its reference morphology, in Table 5.2. Values there are rounded for reading; the calculations retain the full-precision common vector. The corresponding linear auxiliary matrices and mass-normalized vectors are

$$M_0 = \begin{bmatrix} .01806210949 & .002942788035 \\ .002942788035 & .001025996728 \end{bmatrix}, \qquad K_0 = \begin{bmatrix} .6536226658 & -.001187311181 \\ -.001187311181 & .3019920619 \end{bmatrix},$$
$$u_{\mathrm{IP}} = (6.99370047,\ 2.65567809)^{\mathsf{T}}, \quad \Omega_{\mathrm{IP}} = 5.835719714, \tag{D.13}$$
$$u_{\mathrm{AP}} = (7.41765930,\ -42.69238689)^{\mathsf{T}}, \quad \Omega_{\mathrm{AP}} = 24.230934379.$$

The large Euclidean entries reflect normalization in the small effective-inertia metric, not large physical joint excursions.

The activities in (5.19) give $S_j = A_j/(A_{\mathrm{IP}} + A_{\mathrm{AP}})$ and $R_j = A_j/A_{j'}$, $j' \neq j$. For either position or rate signals $a(t), b(t)$, use the time-weighted Pearson correlation

$$\mathrm{corr}_T(a,b) = \frac{\int_0^T (a-\bar{a})(b-\bar{b})\ \mathrm{d}t}{\sqrt{\int_0^T (a-\bar{a})^2\ \mathrm{d}t \int_0^T (b-\bar{b})^2\ \mathrm{d}t}}, \qquad \bar{a} = \frac{1}{T}\int_0^T a\ \mathrm{d}t. \tag{D.14}$$

IP requires both joint-angle and joint-rate correlations to be at least .5; AP requires both to be at most $-.5$. In either case the expected sector must satisfy $S_j \geq .65$ and $R_j \geq 2$. The ratio condition already implies $S_j \geq 2/3$; the share threshold records the implemented test. These measures classify finite-amplitude organization, not genealogy from a particular linear eigenvector.

An auxiliary damping parameter can remove the conservative degeneracy in a continuation formulation. In the normalized quasi-velocity coordinates used here, its effect is explicitly

$$M\dot{\xi} + h + \begin{bmatrix} 0 \\ 0 \\ \nabla U \end{bmatrix} + \mu_{\mathrm{aux}}\xi = 0, \qquad \dot{E} = -\mu_{\mathrm{aux}}\xi^{\mathsf{T}}\xi. \tag{D.15}$$

For any nonstationary periodic solution,

$$0 = E(T) - E(0) = -\mu_{\mathrm{aux}} \int_0^T \xi^{\mathsf{T}}\xi\ \mathrm{d}t \quad \Longrightarrow \quad \mu_{\mathrm{aux}} = 0. \tag{D.16}$$

This exact unfolding identity explains the device without adding physical dissipation to the reported passive cycles. The current phase/energy boundary-value correction instead uses the independent equations (5.16).

D.3 Independent checks on the seventeen primitive cycles

The calculation retains eight AP and nine IP samples. Each earlier trace serves only as a predictor: it is corrected at the prescribed energy and then integrated afresh under (D.11). On normalized time $s \in [0,1]$, collocation solves $z' = Tf(z)$ with five periodicity conditions, $\delta_1(0) = 0$, and the initial energy condition. The requested relative residual is 5×10^{-10}, the boundary tolerance 2.5×10^{-12}, the initial mesh has 481 points, and the maximum mesh has 50,000. The selected crossing has $\sigma_1 > 0$.

Two independent DOP853 integrations (Hairer et al., 1993) use the initial state and corrected period. Nominal and refined settings are respectively

	relative tolerance	absolute tolerance	maximum step
nominal	10^{-11}	10^{-13}	$T/1000$
refined	8×10^{-13}	8×10^{-15}	$T/8000$

(D.17)

Every component of the actual integrated terminal state is compared with the initial state. Component differences are divided by fixed scales $s_i = \max\{1, \mathrm{ptp}(z_i^{\mathrm{anchor}})\}$ from the corresponding reference trajectory. Full-state mismatch and nominal/refined endpoint disagreement must each be below 10^{-8} in this scaling. A separate interval-by-interval exact-field check includes the terminal interval; it does not rely on a visually closed polyline.

Primitive-period testing monitors $\delta_1 = 0, \sigma_1 > 0$ through $1.02T$. The crossing near T must agree with the initial state within 10^{-8}, and earlier positive crossings must not be full-state repetitions. The tested period tolerance is $10^{-8}T$. This establishes the first monitored full section repetition numerically, not an interval-arithmetic minimal-period theorem. Some historical AP exports traversed the primitive loop three times. A predictor for one traversal is evaluated at exactly one third of the reference duration before local correction; taking one third of an integer sample count is not equivalent.

Two phase grids serve different purposes. Geometry and phase-aligned comparisons use 1024 points on $[0, 1)$, without double-counting the initial state. Energy-transfer closure uses 1025 points on $[0, 1]$, including the actual independently integrated state at T. With

$$B_\perp = \max_t E_\perp(t) - \min_t E_\perp(t), \qquad \epsilon_{\mathrm{POE}} = \frac{|E_\perp(T) - E_\perp(0)|}{B_\perp}, \tag{D.18}$$

the transfer span must first exceed the numerical resolution floor. The accepted normalized closure is below 10^{-8}. This checks cyclic recovery relative to the energy actually redistributed, rather than dividing a tiny transfer by a much larger total energy. Total-energy variation is checked separately.

The phase-fixed six-by-six endpoint Jacobian differentiates the five periodicity conditions and initial-energy equation of Equation (5.16) with respect to the five section coordinates and T. It is evaluated by centred finite differences with relative step 10^{-6} and scaled rows and columns. The numerical rank threshold is 10^{-12} times the largest singular value. Recovering the omitted velocity from energy also requires a nonzero derivative (5.17); for this sample its absolute lower bound is 50, with a normalized noncancellation margin at least .99. These safeguards accompany, rather than replace, the independent six-state endpoint comparison.

The IP endpoint problem is appreciably more ill-conditioned than AP; full floating-point rank should not be read as an exact nonsingularity proof. An additional sampled geometric test compares the scaled flow direction with an energy tangent estimated between phase-aligned samples. Its rank-two result concerns cycle–phase geometry and is distinct from the rank-six endpoint system. Neither test establishes connection across uncomputed energies.

Pose is integrated from the identity using (5.2) and compared between nominal and refined solutions. Its dimensionless nonidentity measure scales translation by $l_1 + L_2 + L_3$ and must exceed both 10^{-5} and 100 times its numerical error. All seventeen samples pass with nonzero translation. The additional AP target at energy 220.8783532516 did not converge under the prescribed correction. It remains a failed attempt, not an eighteenth cycle or evidence that a physical family ends there.

D.4 The separate atlas and local design comparisons

The atlas uses 41 AP and 23 IP supports recomputed at primitive period, not the seventeen-sample set above. Across these 64 supports, the maximum scaled independently integrated endpoint mismatch is 2.13×10^{-14}, the maximum collocation residual 5.00×10^{-10}, the maximum relative energy span 3.85×10^{-13}, and the maximum normalized POE closure 2.75×10^{-10}. These checks support the plotted periodic cycles. The sheets and intermediate highlighted cuts remain interpolations.

Thirty-nine supports fail the same absolute lower bound on $|\partial E/\partial v|$ used above. Their independently integrated six-state periodicity nevertheless passes: failure of this stronger speed-recovery safeguard is not failure of internal periodicity. Thus only 25 meet every configured diagnostic; neither the drawing nor these counts transfer the seventeen-sample pose and rank results to the whole atlas.

Table D.1: Numerical checks for the seventeen independently corrected samples only. Values are rounded extrema, not rigorous interval bounds. The rank is that of the phase-fixed endpoint/energy system.

Quantity	AP (8 samples)	IP (9 samples)
Maximum refined scaled state mismatch	9.77×10^{-14}	1.30×10^{-9}
Maximum relative total-energy variation	2.98×10^{-14}	2.41×10^{-14}
Maximum normalized POE closure	8.47×10^{-10}	8.07×10^{-9}
Minimum expected modal share	.701559	.964613
Numerical endpoint-system rank	6	6
Scaled condition-number range	$(4.83\text{–}4.92) \times 10^5$	1.95×10^8–2.07×10^9

The design comparisons use the $G/K/D$ blocks of Section 5.6 and preserve all unselected coefficients. For replay, the stored continuation coordinate is $1 - \alpha$ in (5.21). The four distinct outcomes are given in Table 5.3: in particular, the halfway elastic IP energy balance is not a passive cycle, and unattempted points after a stopped correction have no inferred outcome.

Appendix E

Forced-response equations and numerical verification

This appendix supplies the equations and numerical definitions of the two constructions in Chapter 6. The first section specifies the continuation from a natural cycle, including how the initial torque phase is chosen from its work balance. The remaining sections specify the different mechanical model used for the broad forced-response study, the phase and speed measurements, and the checks on periodicity and transport. Keeping these definitions together distinguishes a model equation, a measured observable and a numerical approximation of a solution.

E.1 Numerical construction of the mechanics-preserving bridge

The bridge retains the mechanical parameters

$$\varepsilon = 0.681, \qquad \gamma = 2.504, \qquad k_2 = 1, \qquad k_4 = 10 \tag{E.1}$$

of Chapter 4. Only the declared viscous and input terms are added:

$$\mu_{\mathrm{visc}}(\lambda) = \lambda c_\star, \qquad u(t;\lambda) = \lambda A_\star \sin\!\left(\frac{2\pi t}{T_\lambda} + \phi_\lambda\right), \qquad c_\star = 10^{-3}. \tag{E.2}$$

The initial torque phase is chosen from the conservative waveform rather than guessed. With $\Omega_0 = 2\pi/T_0$, define its joint-rate correlations and the work that the added loss would remove at first order:

$$\begin{aligned} S_\sigma &= \int_0^{T_0} \sigma(t)\sin(\Omega_0 t)\ \mathrm{d}t, & C_\sigma &= \int_0^{T_0} \sigma(t)\cos(\Omega_0 t)\ \mathrm{d}t, \\ R_\sigma &= \sqrt{S_\sigma^2 + C_\sigma^2}, & W_{\mathrm{loss},0} &= c_\star \int_0^{T_0} (v^2 + \omega^2 + \sigma^2)\ \mathrm{d}t. \end{aligned} \tag{E.3}$$

On this unchanged waveform the first-order input work, after dividing by λ, is $A(S_\sigma \cos\phi + C_\sigma \sin\phi)$. Its largest value is AR_σ, attained at $\phi_{\max} = \operatorname{atan2}(C_\sigma, S_\sigma)$. Thus, for $R_\sigma > 0$, the smallest amplitude able to balance this first-order loss is $A^{(0)}_{\min} = W_{\mathrm{loss},0}/R_\sigma$. For the refined anchor it is 0.03376316109. The computation uses a 10% margin:

$$A_\star = 0.03713947719. \tag{E.4}$$

The equation $A_\star R_\sigma \cos(\phi - \phi_{\max}) = W_{\mathrm{loss},0}$ then has two distinct phase solutions; initialization uses

$$\phi_0 = \phi_{\max} - \arccos\!\left(A^{(0)}_{\min}/A_\star\right). \tag{E.5}$$

The margin avoids tangency at maximum input work. This scalar calculation initializes the continuation; it does not replace correction of the full periodicity equations. At $\lambda = 0$ the input vanishes, so its physical phase remains undefined.

Time zero is fixed by $\delta(0) = 0$, $\sigma(0) > 0$. For every positive λ, the unknowns are $(v_0, \omega_0, \sigma_0, T_\lambda, \phi_\lambda)$; the equations are the four synchronized internal-periodicity equations and the time-weighted mean-speed

condition $\bar{v} = 10.435810958679$. The drive frequency $2\pi/T_\lambda$ therefore changes with the corrected response. Nineteen prescribed homotopy values are used:

$$\lambda \in \{0, 10^{-4}, 3\times10^{-4}, 10^{-3}, 3\times10^{-3}, 5\times10^{-3}, 10^{-2}, 3\times10^{-2}, 0.06, 0.1, 0.2, \dots, 1\}. \tag{E.6}$$

Every declared point was reached. The maximum strict reduced-return residual is 7.102×10^{-10}, the maximum mean-speed error is 1.075×10^{-12}, and the maximum absolute defect in the differential energy identity is 4.112×10^{-13}. The largest relative cycle work imbalance is 2.940×10^{-4} at $\lambda = 10^{-4}$, where both work terms are approximately 10^{-5}; at $\lambda = 1$ the relative imbalance is 2.60×10^{-11}. Meshes from 501 to 4001 points and three independent integration tolerances were compared. The largest strict–ultra endpoint difference is 5.405×10^{-14}.

The waveform comparison uses $N = 4001$ equally spaced normalized times $s_k \in [0, 1]$, including both endpoints, and $Z_{\lambda,k} = z_\lambda(T_\lambda s_k)$. The common oriented section aligns the cycles; no phase shift is fitted. With anchor scales $s_\delta = \mathrm{ptp}(\delta_0)$, $s_v = \mathrm{RMS}(v_0 - \bar{v})$, $s_\omega = \mathrm{RMS}(\omega_0)$ and $s_\sigma = \mathrm{RMS}(\sigma_0)$ on that grid, where RMS is root mean square and ptp is maximum minus minimum, the distance is

$$d_\lambda = \left[\frac{1}{4N}\sum_{k=1}^{N}\sum_{j=1}^{4}\left(\frac{Z_{\lambda,k,j} - Z_{0,k,j}}{s_j}\right)^2\right]^{1/2}. \tag{E.7}$$

Each scale is bounded below by 10^{-12} in the computation. Scaling v by its fluctuation, not its large mean, keeps changes in the propulsion waveform visible. The distance is 4.658×10^{-5} at $\lambda = 10^{-4}$ and 0.28708 at $\lambda = 1$: the computed weakly forced waveform approaches this anchor, with progressively larger deformation along this one branch.

E.2 Equations of the parameter-distinct forced model

Let $\eta_f = (v, \omega_b, \sigma)^\mathsf{T}$. The response computation uses

$$M_f(\delta) = \begin{bmatrix} 1 & -(\varepsilon_f - 1)^2 \sin\delta & -(\varepsilon_f - 1)^2 \sin\delta \\ -(\varepsilon_f - 1)^2 \sin\delta & M_{22} & M_{23} \\ -(\varepsilon_f - 1)^2 \sin\delta & M_{23} & M_{33} \end{bmatrix}, \tag{E.8}$$

where

$$\begin{aligned} M_{22} &= \varepsilon_f^3\gamma_f + \varepsilon_f^3 - \gamma_f(\varepsilon_f - 1)^3 \\ &\quad + (1 - \varepsilon_f)\left[-2\varepsilon_f^2\cos\delta + 2\varepsilon_f^2 + 2\varepsilon_f\cos\delta - 2\varepsilon_f + 1\right], \end{aligned} \tag{E.9}$$

$$M_{23} = (\varepsilon_f - 1)^2\left[-\varepsilon_f\gamma_f + \varepsilon_f\cos\delta - \varepsilon_f + \gamma_f + 1\right], \tag{E.10}$$

$$M_{33} = -(\varepsilon_f - 1)^3(\gamma_f + 1). \tag{E.11}$$

The continued data use $\varepsilon_f = 0.7317084397592031$ and $\gamma_f = 0.736801492464307$. A numerical minimization over one full angular period, $\delta \in [-\pi, \pi]$, gives minimum eigenvalue 0.02561150819. This supports positive definiteness of the inertia over that interval to the reported numerical precision; it is not an interval-certified lower bound.

The code evaluates the acceleration in right-hand-side form,

$$M_f(\delta)\dot{\eta}_f = b_f(\delta, \eta_f) + B_\delta\tau_\delta - c_d\eta_f, \qquad B_\delta = \begin{bmatrix}0 & 0 & 1\end{bmatrix}^\mathsf{T}, \tag{E.12}$$

with

$$b_f = \begin{bmatrix} \varepsilon_f\omega_b^2 + (1 - \varepsilon_f)^2(\sigma + \omega_b)^2\cos\delta \\ \varepsilon_f(1 - \varepsilon_f)^2(\sigma + 2\omega_b)\sigma\sin\delta - \omega_b v\left[\varepsilon_f + (1 - \varepsilon_f)^2\cos\delta\right] \\ -\varepsilon_f(1 - \varepsilon_f)^2\omega_b^2\sin\delta - (1 - \varepsilon_f)^2\omega_b v\cos\delta - 2\delta \end{bmatrix}. \tag{E.13}$$

Thus the left-hand-side convention of Equation (6.9) is exact when

$$F_f = -b_f, \qquad D_f = c_d I_3. \tag{E.14}$$

In particular, the third component of F_f ends in $+2\delta$. The factor two follows from the implemented potential $V_f(\delta) = \delta^2$, for which the spring torque is -2δ. Writing the vector b_f on the left with a plus sign would

reverse every conservative acceleration term and is not the implemented model. In the same convention, the stored energy is $E_f = \eta_f^{\mathsf{T}} M_f(\delta)\eta_f/2 + \delta^2$. Differentiating it along the field gives $\dot{E}_f = \tau_\delta \sigma - c_d \eta_f^{\mathsf{T}} \eta_f$, which is the sign-sensitive identity used in the work check below.

To express the sinusoidal input in autonomous form, introduce the Cartesian forcing-clock state (a, b). Its dynamics and torque output are

$$\dot{a} = -\Omega b + a(1 - a^2 - b^2), \qquad \dot{b} = \Omega a + b(1 - a^2 - b^2), \qquad \tau_\delta = A_{\mathrm{exc}} b. \tag{E.15}$$

Writing $r^2 = a^2 + b^2$ gives $\dot{r} = r(1 - r^2)$ and, for $r > 0$, $\dot{\psi} = \Omega$. Thus $r = 1$ is an invariant unit circle on which $a = \cos\psi$, $b = \sin\psi$; the added radial term maintains the numerical clock and is not part of the robot's mechanical dissipation. The implementation uses $A_{\mathrm{exc}} b$ directly; it does not divide by $\sqrt{a^2 + b^2}$.

E.2.1 Checking the stored responses against the equations of motion

For all 13,782 continued cycles, the first and last recorded values agree for δ, v, ω_b, σ, a, and b at the written precision. Endpoint agreement alone could also occur in an incorrectly exported closed curve. Dynamic consistency is therefore tested separately. On each of the 2,756,400 stored intervals, including cycle closure, two fourth-order Runge–Kutta half steps of Equations (E.12) and (E.15) predict $\widehat{z}_{k+1}$ from z_k. The scaled defect divides coordinate i by $s_i = \max\{1, \mathrm{ptp}(z_i)\}$ on that cycle. For cycle j, the continued forcing frequency used by the field is inferred from the recorded clock turn as $\Omega_j = [\mathrm{unwrap}\,\psi_j(T_j) - \mathrm{unwrap}\,\psi_j(0)]/T_j$, with $\psi_j = \mathrm{atan2}(b_j, a_j)$, rather than read from a single historical archive scalar. The population maxima are listed in Table E.1. They are mesh-dependent local multiple-shooting defects, not a mesh-independent norm of field error; the interval count, Richardson estimate, and DOP853 comparison state their scale.

The one-step/two-half-step Richardson estimate is at most 9.73×10^{-7}. Tight-tolerance DOP853 integration (Hairer et al., 1993, Section II.5) of the worst raw interval gives a defect 4.837×10^{-4} and differs from the two-half-step endpoint by 9.38×10^{-7}. The independently coded numerical field agrees with the symbolic model builder to 1.42×10^{-13} over 72 state–parameter evaluations. Finally, Simpson integration of every cycle verifies

$$\int_0^T A_{\mathrm{exc}} b\,\sigma \,\mathrm{d}t = \int_0^T c_d(v^2 + \omega_b^2 + \sigma^2)\,\mathrm{d}t \tag{E.16}$$

with maximum relative imbalance 7.583×10^{-6} and median 5.257×10^{-9}.

Unwrapping $\psi = \mathrm{atan2}(b, a)$ gives one forcing-clock turn per mechanical return for every cycle, with maximum winding-number deviation 2.22×10^{-16} and maximum clock-radius defect 2.86×10^{-9}. Together, recorded closure and the local-flow audit support the primary 1:1 synchronized interpretation at the reported numerical resolution. They do not establish attraction, reachability, or global regularity of the solution set. The unit winding expresses synchronization, not resonance by itself: *resonant* refers to the component seeded near the small-amplitude natural period and followed through the nonlinear response ridge; response maxima and quadrature crossings are assessed separately.

All variables in this computation are nondimensional. For reference mass M_0, length l_0, and the coefficient k_0 in the dimensional potential $V = k_0\delta^2$, the time scale and torque-amplitude conversion are

$$t_c = \sqrt{\frac{M_0 l_0^2}{k_0}}, \qquad A_{\mathrm{exc}} = \frac{\tau_{\mathrm{amp}} t_c^2}{M_0 l_0^2} = \frac{\tau_{\mathrm{amp}}}{k_0}. \tag{E.17}$$

Consequently, $\tau_{\mathrm{amp}} = k_0 A_{\mathrm{exc}}$. The numerical color coordinate used in the response plots equals $1000 A_{\mathrm{exc}}$; it is an equivalent N mm value only for $k_0 = 1\,\mathrm{N\,m}$.

Table E.1: Seam-inclusive local-flow defects over the complete forced-response population. The closure column concerns the last stored interval returning to the initial state.

State coordinate	maximum raw defect	maximum scaled defect	maximum raw closure defect
δ	3.302×10^{-6}	2.680×10^{-7}	4.391×10^{-7}
v	1.660×10^{-4}	1.716×10^{-5}	6.876×10^{-5}
ω_b	1.471×10^{-4}	1.175×10^{-5}	6.679×10^{-5}
σ	4.827×10^{-4}	3.331×10^{-6}	1.305×10^{-4}
a	1.356×10^{-8}	6.781×10^{-9}	1.356×10^{-8}
b	2.888×10^{-8}	1.444×10^{-8}	2.799×10^{-8}

The 25 fixed nondimensional forcing-amplitude levels, reported to eight decimal places and in increasing order, are

$$\begin{aligned}\mathcal{A} = \{&0.00204965,\ 0.00329097,\ 0.00451922,\ 0.00577260,\ 0.00712466,\\ &0.00865827,\ 0.01043898,\ 0.01251222,\ 0.01490514,\ 0.01762767,\\ &0.02067394,\ 0.02401735,\ 0.02763993,\ 0.03148703,\ 0.03551893,\\ &0.03969468,\ 0.04394984,\ 0.04821334,\ 0.05243581,\ 0.05655617,\\ &0.06052866,\ 0.06426299,\ 0.06771534,\ 0.07084794,\ 0.07357499\}.\end{aligned} \tag{E.18}$$

An integer subscript used below only orders these physical amplitude values; it is not an additional discrete model parameter.

E.3 Phase extracted from the continued cycles

The response map stores the phase of joint deformation relative to applied joint torque. The continuation software AUTO-07P, accessed through the PyDSTool/PyCont interface (Doedel et al., 2007; Clewley, 2012), produces a mesh that is not uniform in time. Each cycle is first interpolated periodically onto 1,000 uniformly spaced times in $[0, T)$; the duplicated endpoint is excluded. For a resampled signal y, the extractor removes its mean and forms the analytic signal

$$a_y(t) = y(t) - \langle y \rangle + \mathrm{i}\,\mathcal{H}[y - \langle y \rangle](t), \tag{E.19}$$

where $\mathcal{H}$ denotes the discrete Hilbert transform. The cycle phase is

$$\phi_{\delta\tau} = \operatorname{Arg}\left[\frac{1}{N}\sum_{k=1}^{N}\exp\{\mathrm{i}\big(\arg a_\delta(t_k) - \arg a_\tau(t_k)\big)\}\right]. \tag{E.20}$$

Thus the signal order is joint angle minus applied torque, the wrap interval is $[-\pi, \pi)$, and quadrature is $\phi_{\delta\tau} = -\pi/2$. This is a Hilbert analytic-signal statistic. A Fourier-coefficient phase can approximate it on a nearly harmonic response, but the two definitions are not identical on a distorted waveform. On the same centered $N = 1000$ samples, let $\widehat{\delta}_k$ denote the discrete Fourier coefficients. The plotted total harmonic distortion is

$$\mathrm{THD}_\delta = \frac{\sqrt{\sum_{k=2}^{N/2}|\widehat{\delta}_k|^2}}{|\widehat{\delta}_1|}. \tag{E.21}$$

It is a dimensionless ratio, not a percentage: zero describes a purely fundamental waveform, and larger values mean more higher-harmonic content relative to the fundamental. Both phase and THD require a resolved nonzero oscillatory signal; neither is a reliable timing diagnostic at rest.

E.4 Peak-speed and quadrature extraction

For each amplitude level $A_r \in \mathcal{A}$, let

$$v_{\mathrm{pk},rj} = \max_k v_{rj}(t_k), \qquad j_{p,r} \in \arg\max_j v_{\mathrm{pk},rj}, \qquad (T_{p,r}, v_{p,r}) = (T_{rj_{p,r}}, v_{\mathrm{pk},rj_{p,r}}) \tag{E.22}$$

where k indexes the recorded within-cycle samples and j follows continuation order. Thus the stated peak is a sampled positive maximum, not a spline estimate between nodes. An exact sampled $-\pi/2$ value is used directly. Otherwise, a quadrature crossing requires two consecutive responses bracketing that target. If

$$d_j = \phi_{rj} + \frac{\pi}{2}, \qquad d_j d_{j+1} < 0, \tag{E.23}$$

the interpolation fraction and crossing speed are

$$\begin{aligned}\lambda_j &= \frac{|d_j|}{|d_j| + |d_{j+1}|},\\ T_{q,r} &= (1-\lambda_j)T_{rj} + \lambda_j T_{r,j+1},\\ v_{q,r} &= (1-\lambda_j)v_{\mathrm{pk},rj} + \lambda_j v_{\mathrm{pk},r,j+1}.\end{aligned} \tag{E.24}$$

When several crossings occur, the first in continuation order is used. No crossing is inferred when the target is absent.

The reported proximity metric is

$$\epsilon_{v,r} = 100\frac{|v_{q,r} - v_{p,r}|}{|v_{p,r}|}, \tag{E.25}$$

with the corresponding period gap defined analogously. Both numerical values are reported in percent. Table E.2 separates the three observed forcing-amplitude regimes.

Table E.2: Quadrature verification by forcing-amplitude interval.

Forcing-amplitude levels	Crossing status	Relation to the observed speed peak
$0.00204965 \le A_{\mathrm{exc}} \le 0.03551893$ (15 levels)	all 15 cross $-\pi/2$	maximum linearly interpolated speed gap 0.0635625%; maximum period gap 0.0289515%
$0.03969468 \le A_{\mathrm{exc}} \le 0.05655617$ (five levels)	all five cross	crossing belongs to a much slower part of the continued response, 59.19–71.84% below the observed peak
$0.06052866 \le A_{\mathrm{exc}} \le 0.07357499$ (five levels)	no sampled crossing	the nearest recorded value is not promoted to a quadrature point

The main text rounds the lowest-15-level bound upward to 0.064%. The largest phase error evaluated at an observed speed peak on those levels is 0.0452127 rad. Near-equality of peak speed and quadrature speed is therefore a local property of the rising branch, not a property of all 25 forcing-amplitude levels.

E.5 Population-wide pose reconstruction

The continued-cycle data contain the time histories needed to reconstruct pose for every response. With $\theta(0) = 0$, the calculation uses

$$\theta(t) = \int_0^t \omega_b(s)\ \mathrm{d}s, \qquad \Delta x = \int_0^T v(t)\cos\theta(t)\ \mathrm{d}t, \qquad \Delta y = \int_0^T v(t)\sin\theta(t)\ \mathrm{d}t. \tag{E.26}$$

Changing the initial pose rigidly transforms the reconstructed path and cannot change whether Δg is the identity.

Figure E.1 illustrates this integration for three peak-speed responses from the same forced population. The vertical offsets separate drawings, not initial conditions used to compute their translation. The different path lengths are outputs of the selected cycles, not a ranking of locomotor efficiency or equal-time performance.

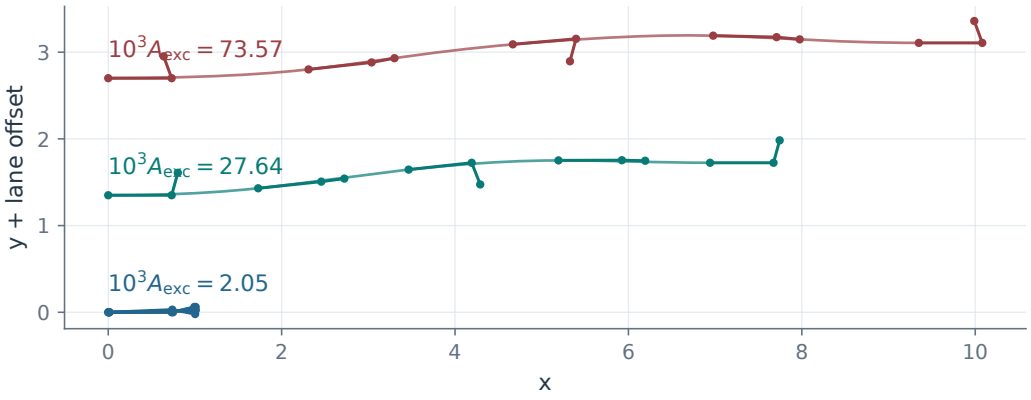


Figure E.1: Pose reconstructions for peak-speed responses at $A_{\mathrm{exc}} = 0.00204965$, 0.02763993, and 0.07357499. Artificial vertical offsets separate the paths. Each displayed reduced response returns over one forcing period while the reconstructed carrier advances.

Composite trapezoidal integration on the nonuniform continuation mesh gives a nonzero translation for all 13,782 responses. Repeating the reconstruction with cumulative Simpson integration for θ and Simpson

integration for $(\Delta x, \Delta y)$ gives

$$3.46365 \times 10^{-7} \le \sqrt{\Delta x^2 + \Delta y^2} \le 9.35730, \qquad \max \frac{\|\Delta p_{\mathrm{S}} - \Delta p_{\mathrm{T}}\|}{\|\Delta p_{\mathrm{S}}\|} = 0.6776\%, \tag{E.27}$$

where $\Delta p = (\Delta x, \Delta y)$ and subscripts S and T denote the Simpson and trapezoidal reconstructions. In particular, even the smallest reconstructed translation remains nonzero under both rules. Their agreement is a numerical consistency test, not a rigorous integration-error bound for such a small displacement. Together with the local-flow checks above, these computations support synchronized internal motion with transport at the reported numerical resolution. The regularity required to interpret neighboring responses as an RLM chart is a separate condition, stated in Section 6.4; pose quadrature does not test it.

E.6 Nonlinear shift from the small-amplitude period

The reference period is that of the conservative linearization of this forced model's mechanics, not the maximum of its damped speed-response curve. At $\delta = v = \omega_b = \sigma = 0$, the yaw and joint equations without loss or input are $M_{22}(0)\dot{\omega}_b + M_{23}(0)\ddot{\delta} = 0$ and $M_{23}(0)\dot{\omega}_b + M_{33}(0)\ddot{\delta} + 2\delta = 0$. Eliminating $\dot{\omega}_b$ gives $M_{\mathrm{lin}}\ddot{\delta} + 2\delta = 0$, with

$$M_{\mathrm{lin}} = M_{33}(0) - \frac{M_{23}(0)^2}{M_{22}(0)} = 0.0258224048916. \tag{E.28}$$

Since the linear restoring coefficient is 2,

$$T_{\mathrm{lin}} = 2\pi\sqrt{M_{\mathrm{lin}}/2} = 0.7139424640185. \tag{E.29}$$

The nearby value 0.7139425323 used in an earlier response calculation is the period of a finite seed cycle, not the result of this linearization. At the observed speed maximum for $A_{\mathrm{exc}} = 0.03551893$, $T_p(0.03551893) = 0.7410295954$, giving

$$100\frac{T_p(0.03551893) - T_{\mathrm{lin}}}{T_{\mathrm{lin}}} = 3.7940\%. \tag{E.30}$$

At $A_{\mathrm{exc}} = 0.07357499$ the observed peak period is 0.7411539715, a 3.8114% shift. The period shift depends only on the recorded periods and the mechanical linearization; it is independent of the Hilbert-phase and THD extraction rules. Together, these audits support the declared primary 1:1 synchronized, transport-bearing population at the reported numerical resolution; they do not establish uniqueness, attraction, global regularity, or experimental robustness.

Appendix F

Phase–speed controller and sampled finite-harmonic evidence

The controller of Chapter 7 acts on two measured quantities, but reproducing its simulations also requires the estimator windows, sampled updates, bounds and stage protocol. These are specified here in the order in which information passes through the implementation: signals become phase and mean-speed estimates; these estimates adjust frequency and amplitude; a separate numerical supervisor evaluates completed stages. The final sections compare whole controlled cycles and describe the finite-harmonic models used to assess local feedback behavior.

The feedback's rolling windows are distinct from the longer terminal window used to evaluate a stage. Likewise, the simulator's ability to restart from a saved state is distinct from a feedback recovery action on a physical robot. These two time scales correspond to the solid feedback paths and the separate dashed supervision path in Figure 7.3. The numerical definitions below make that visual separation precise.

F.1 Rolling phase and speed estimates

Let $t_n = n\Delta t$, with $\Delta t = 0.006$, let $T_0 = 2\pi/\Omega_0$ be the period associated with the stage's initial forcing frequency, and let $T_n = 2\pi/|\Omega_{n-1}|$, for $n \geq 1$, be the period estimate before the current controller update. The phase fit uses the last

$$N_{\phi,n} = \max\left\{24, \left\lfloor \frac{0.25T_0}{\Delta t} \right\rfloor, \left\lfloor \frac{T_n}{\Delta t} \right\rfloor\right\} \tag{F.1}$$

samples. Over the frequencies reached here, the active term is one current forcing period. The speed estimate uses the longer window

$$N_{v,n} = \max\left\{N_{\phi,n}, \left\lfloor \frac{3T_n}{\Delta t} \right\rfloor\right\}, \qquad \bar{v}_n = \frac{1}{N_{v,n}} \sum_{j=n-N_{v,n}+1}^{n} v_j. \tag{F.2}$$

Thus phase is updated from approximately one period and mean speed from approximately three periods. Here every v_j is the longitudinal velocity component resolved in the body frame. Consequently $\bar{v}_n$ is the local speed coordinate supplied to the controller; it is not generally the norm of the net inertial translation divided by elapsed time when yaw varies.

For a signal y_i sampled at excitation phases φ_i on the phase window, define

$$X_{K_{\mathrm{est}}} = \begin{bmatrix} 1 & \sin\varphi_1 & \cos\varphi_1 & \cdots & \sin K_{\mathrm{est}}\varphi_1 & \cos K_{\mathrm{est}}\varphi_1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & \sin\varphi_N & \cos\varphi_N & \cdots & \sin K_{\mathrm{est}}\varphi_N & \cos K_{\mathrm{est}}\varphi_N \end{bmatrix}. \tag{F.3}$$

The coefficients are obtained by a linear least-squares solve,

$$\widehat{\beta}_y = X_{K_{\mathrm{est}}}^{\dagger}\boldsymbol{y} = \begin{bmatrix} \widehat{c}_{y,0} & \widehat{a}_{y,1} & \widehat{b}_{y,1} & \cdots \end{bmatrix}^{\mathsf{T}}. \tag{F.4}$$

The dagger denotes the least-squares solution; the normal-equation inverse is not formed explicitly. With $K_{\mathrm{est}} = 3$, the fundamental phase is

$$\phi_y = \operatorname{atan2}(\widehat{b}_{y,1}, \widehat{a}_{y,1}), \qquad e_{\phi,n} = \operatorname{wrap}_{[-\pi,\pi)}\left(\phi_\delta - \phi_\tau + \frac{\pi}{2}\right). \tag{F.5}$$

The same excitation clock and design matrix are used for δ and τ_δ. Their order is response minus torque. Fit amplitude, residual RMS, and the condition number of $X_{K_{\mathrm{est}}}$ diagnose a degenerate fundamental or an inadequately covered window; only the fundamental phase difference enters feedback. Indeed, $\widehat{a}_{y,1}\sin\varphi + \widehat{b}_{y,1}\cos\varphi = \widehat{R}_y\sin(\varphi+\phi_y)$, which explains the order of the two arguments in $\operatorname{atan2}$. The known accumulated excitation phase labels the samples; the fit does not assume a constant frequency throughout a transient window.

F.2 Sampled phase and amplitude laws

For a scalar r, write $\operatorname{sat}_{[\ell,u]}(r) = \min\{u, \max\{\ell, r\}\}$. Once the phase estimate is available, the integral state and the unsaturated position form of the phase PI are

$$I_{\phi,n} = I_{\phi,n-1} + \Delta t\, e_{\phi,n}, \qquad \Omega_n^{\mathrm{raw}} = \Omega_0 + s\,(K_P e_{\phi,n} + K_I I_{\phi,n}). \tag{F.6}$$

The actual sampled update uses the equivalent velocity form and applies the frequency bound at every step:

$$\Omega_n = \operatorname{sat}_{[5,11]}(\Omega_{n-1} + s\,[K_P(e_{\phi,n} - e_{\phi,n-1}) + K_I\Delta t\, e_{\phi,n}]). \tag{F.7}$$

Away from saturation, summing Equation (F.7) gives Equation (F.6). The previous phase error is initialized to zero, and Ω_0 is the excitation frequency carried into the stage (the Ω_{free} of the condensed position form). For the output-minus-input convention in Equation (F.5), the locally verified sign is $s = +1$; the gains are $K_P = 0.04$ and $K_I = 0.08$. The saturation limits in Equation (F.7) are frequencies per model-time unit.

The amplitude update starts only after both a phase estimate and the three-period speed estimate are available. Its phase-health weight is

$$g_{\phi,n} = \begin{cases} 1, & |e_{\phi,n}| \le 0.20, \\ (0.55 - |e_{\phi,n}|)/0.35, & 0.20 < |e_{\phi,n}| < 0.55, \\ 0, & |e_{\phi,n}| \ge 0.55. \end{cases} \tag{F.8}$$

With $e_{v,n} = \bar{v}_{\mathrm{des}} - \bar{v}_n$, the nominal commanded amplitude rate and bounded update are

$$u_{A,n} = \operatorname{sat}_{[-r_A,r_A]}\left(g_{\phi,n}K_v e_{v,n} - \mathbf{1}_{\{g_{\phi,n}=0\}}\rho(A_{\mathrm{exc},n-1} - A_{\min})\right), \tag{F.9}$$

$$A_{\mathrm{exc},n} = \operatorname{sat}_{[A_{\min},A_{\max}]}(A_{\mathrm{exc},n-1} + \Delta t\, u_{A,n}). \tag{F.10}$$

The values are

$$K_v = 2.5\times10^{-5}, \qquad r_A = 2.5\times10^{-4}, \qquad \rho = 10^{-4}, \qquad (A_{\min}, A_{\max}) = (10^{-4}, 8\times10^{-2}). \tag{F.11}$$

The middle branch of g_ϕ progressively slows the speed adaptation; full loss of phase lock closes the speed channel and adds a weak decay toward $A_{\min}$. As in Chapter 6, forcing amplitudes are nondimensional, and r_A is the corresponding amplitude rate per model-time unit. At each sample the signals are recorded and the two updates are evaluated; the new amplitude and frequency are then held over the next integration step, while the excitation phase advances continuously at that frequency. Before their respective windows are filled, frequency and amplitude retain their stage-initial values. Estimator fit diagnostics are recorded, but no separate fit-quality rejection threshold is implemented in this nominal protocol.

F.3 Local sign tests and the response segment

The responses used for interpretation are ordered by increasing A_{exc}. Table F.1 lists the values that delimit the controller tests and the sampled finite-harmonic analysis. The ordering is not supplied to the online controller.

The final column uses the stored offline Hilbert phase, as in the forced-response figures. It identifies the reference responses; the local sign tests below instead evaluate the online fundamental-fit phase error.

Table F.1: Representative ordered forcing-amplitude levels.

Role	A_{exc}	$\bar{v}$	Ω_{exc}	$\phi + \pi/2$ [rad]
lowest stored response	0.002050	0.0215	8.8006	0.00237
low interior response	0.007125	1.7977	8.7948	-0.01262
phase-sign calibration	0.017628	6.2112	8.7439	0.000024
last linearized level	0.031487	10.8764	8.5987	0.04259
next numerical endpoint	0.035519	12.3649	8.4790	0.05127

The preliminary sign tests last 150 reference periods with step 0.002. They use $(K_P, K_I) = (0.02, 0.05)$ for the $\pm 0.5\%$ detunings and reversed sign, and $(0.05, 0.05)$ for the -3% detuning, with frequency bounds $[0.1, 50]$. These settings differ from the subsequent nominal two-loop test. For their reported tail-refit statistic, a single three-harmonic fit uses the final quarter of the run and the absolute error from $-\pi/2$ is recorded.

At the phase-sign calibration level, a -0.5% initial frequency detuning gave absolute tail-refitted phase error $2.391 \times 10^{-3}\,\mathrm{rad}$ with $s = +1$, while $s = -1$ gave $1.2452\,\mathrm{rad}$. Correct-sign starts at $+0.5\%$ and -3% detuning ended at 2.679×10^{-3} and $4.103 \times 10^{-3}\,\mathrm{rad}$. The corresponding RMS values of the online rolling error are respectively 0.02034, 0.02059 and $0.03015\,\mathrm{rad}$ for the three correct-sign detunings, and $1.75816\,\mathrm{rad}$ for the reversed sign. Both measures support the local sign choice, but they summarize different windows. By contrast, the 15-run amplitude staircase kept small phase errors but remained on the low response sheet. These two tests distinguish local phase locking from response selection.

F.4 Stage-transition logic and terminal metrics

Each proposed speed stage lasts 150 model-time units. Acceptance is evaluated on its last quarter, $\mathcal{W}_{\mathrm{tail}} = [112.5, 150]$, rather than on either rolling feedback window. For any signal y on a window $\mathcal{W}$ of duration $|\mathcal{W}|$, define

$$\langle y \rangle_{\mathcal{W}} = \frac{1}{|\mathcal{W}|} \int_{\mathcal{W}} y(t)\, \mathrm{d}t, \quad \mathrm{RMS}_{\mathcal{W}}(y) = \sqrt{\langle y^2 \rangle_{\mathcal{W}}}, \quad \mathrm{ptp}_{\mathcal{W}}(y) = \max_{\mathcal{W}} y - \min_{\mathcal{W}} y. \tag{F.12}$$

The numerical tests use the corresponding sample mean, root mean square, and sample maximum minus minimum. Set $\bar{v}_{\mathrm{new}} = \langle v \rangle_{\mathcal{W}_{\mathrm{tail}}}$. A proposal passes the response-health tests when

$$\begin{aligned} \mathrm{RMS}_{\mathcal{W}_{\mathrm{tail}}}(e_\phi) &\le 0.12, \\ \mathrm{ptp}_{\mathcal{W}_{\mathrm{tail}}}(\delta) &\le 8.5, \\ |\langle \omega \rangle_{\mathcal{W}_{\mathrm{tail}}}| &\le 0.08, \\ \bar{v}_{\mathrm{new}} &\ge \bar{v}_{\mathrm{old}} - \max(0.75, 0.12|\bar{v}_{\mathrm{old}}|). \end{aligned} \tag{F.13}$$

Here ω is body yaw rate, not excitation frequency. For a requested speed $\bar{v}_{\mathrm{des}} \ge 8$, acceptance also requires

$$|\bar{v}_{\mathrm{new}} - \bar{v}_{\mathrm{des}}| \le \max(0.65, 0.14|\bar{v}_{\mathrm{des}}|). \tag{F.14}$$

The first proposed speed increment is 2, with nominal limits $0.125 \le \Delta\bar{v}_{\mathrm{des}} \le 2$. After a speed-converged accepted stage it is multiplied by 1.15; after an accepted but off-command stage below the strict high-speed threshold it is multiplied by 0.70. A rejected stage restores the previous accepted plant state, excitation phase, amplitude, and frequency, and halves the increment. Once the accepted target is at least 8, the next increment is capped at 0.25. An initialization terminates after 40 proposals or when a rejected increment falls below 0.125. The speed-converged classification uses Equation (F.14) at every target; below 8 it changes the next step size, but convergence is not required for acceptance. Every new trial clears the rolling histories and initializes the previous phase error to zero. Only the accepted internal state, excitation phase, amplitude and frequency are carried forward. No plant state is physically reversed: rejection means discarding a simulated trial and restarting from a saved state.

Table F.2 summarizes the two realized phase–speed progressions and their terminal forcing amplitudes.

Table F.2: Guarded phase–speed progression from the two tested initializations.

Quantity	Start from rest	Start from the stored lowest-level response
proposed stages	21	21
accepted / rejected	21/0	21/0
terminal mean speed	12.0680391251	12.0677941731
terminal phase-error RMS [rad]	$1.11897244 \times 10^{-3}$	$1.10605507 \times 10^{-3}$
terminal nondimensional A_{exc}	0.0364166157	0.0364087749
A_{exc} of closest stored response by mean speed	0.0355189	0.0355189
online use of the response family	none	none

Across the 42 accepted stages, the maximum phase-error RMS was $2.7012 \times 10^{-3}\,\mathrm{rad}$, the maximum deformation peak-to-peak was 3.5383, and the maximum absolute mean yaw rate was 4.5720×10^{-3}. Every stage therefore remained inside the declared thresholds. Because no proposal was rejected, this nominal campaign does not exercise rejection handling; the separate stress audit below constructs one rejected and one reduced proposal outside the production stage-manager branch.

F.5 Comparing complete cycles and refining the integration step

The terminal speed and phase statistics do not by themselves identify a full response cycle. For a stronger a posteriori comparison, let $z_{\mathrm{c}}(\varphi)$ be the final controlled cycle and $z_r^{(k)}(\varphi)$ the stored peak-response cycle of amplitude-level index k, both sampled at the same $N = 512$ forcing phases without a free cyclic shift. Here $z = (\delta, v, \omega, \sigma)$. With component scales

$$s_j = \max_{k,\ell} z_{r,j}^{(k)}(\varphi_\ell) - \min_{k,\ell} z_{r,j}^{(k)}(\varphi_\ell), \qquad d_k = \left[\frac{1}{4N} \sum_{\ell=1}^{N} \sum_{j=1}^{4} \left(\frac{z_{\mathrm{c},j}(\varphi_\ell) - z_{r,j}^{(k)}(\varphi_\ell)}{s_j} \right)^2 \right]^{1/2}, \tag{F.15}$$

the nearest of the 25 available peak cycles is amplitude level 15. The same calculation also compares the last two controlled cycles and the two endpoints of the last cycle. These are neighborhood and near-periodicity diagnostics, not collocation residuals or membership tests for the complete synchronized- response zero set.

Table F.3 reports the nominal integration step and a replay of the final lowest-amplitude-start stage from exactly the same plant and forcing state, over the same duration, with the step halved.

Table F.3: Terminal full-cycle and time-step audit. All cycle distances are normalized according to Equation (F.15).

Quantity	$\Delta t = 0.006$	$\Delta t = 0.003$
terminal mean speed	12.0677942	12.0677663
terminal phase-error RMS [rad]	1.1060551×10^{-3}	1.1060157×10^{-3}
nearest peak-response amplitude level	15	15
nearest-cycle distance d_{15}	1.7542932×10^{-2}	1.7563618×10^{-2}
last-two-cycle RMS difference	1.1275073×10^{-4}	3.7601361×10^{-5}
one-cycle endpoint ℓ_∞ defect	2.3416024×10^{-4}	5.9930842×10^{-5}

The two complete terminal profiles differ by a normalized RMS of 1.0740512×10^{-4}. Halving the step changes mean speed by 2.7882×10^{-5} and phase-error RMS by $3.94 \times 10^{-8}\,\mathrm{rad}$. The plant was integrated here by explicit fourth-order Runge–Kutta. This single-stage refinement supports the reported terminal observables; it does not cover every stage, the rest-start run, or a controller theorem.

Finally, a stress harness starts from the accepted stage at requested speed 8.2935 and mean speed 7.5801. It deliberately bypasses the nominal high-speed step cap: a direct proposal to 12 ends at mean speed 8.9980, with speed error 3.0020, and is rejected. The rejected terminal state is discarded; the harness explicitly reuses a copy of the accepted plant and forcing state, halves the proposed increment, and accepts target 10.1468 at mean speed 8.7457, with speed error 1.4011. Thus the production simulator and guard are exercised in one constructed reject–reuse–reduced-proposal sequence, but the rejection branch of the production stage manager is not. This is neither a recovery theorem, a basin estimate, nor a test of arbitrary failures.

F.6 Closed finite-harmonic controller matrices

To examine how feedback acts on a nearby waveform, let $z_*(\varphi)$ be one stored response at (A_*, Ω_*). The calculation retains that reference and represents its perturbation with $K_{\mathrm{G}} = 5$ harmonics:

$$z(\varphi) = z_*(\varphi) + \Delta\bar{z} + \sum_{k=1}^{K_{\mathrm{G}}} \big[a_k \sin(k\varphi) + b_k \cos(k\varphi)\big]. \tag{F.16}$$

There are $4(1 + 2K_{\mathrm{G}}) = 44$ coefficient perturbations $x \in \mathbb{R}^{44}$. Write their basis as $B(\varphi)$ and let the coefficients vary with time: $z(t) = z_*(\varphi(t)) + B(\varphi(t))x(t)$. Since $\dot{\varphi} = \Omega$, substitution into $\dot{z} = f(z, A, \varphi)$ gives $B\dot{x} = f(z_* + Bx, A, \varphi) - \Omega(z_*' + B'x)$. Its projection onto B, using 512 uniform forcing phases and the basis Gram matrix, defines $\dot{x} = F_G(x, \Omega, A)$. Centered finite differences then give the local input–output expansion

$$\begin{aligned} \dot{x} &= r_G + A_G x + p_G \Delta\Omega + q_G a + O(\|(x, \Delta\Omega, a)\|^2), \\ \Delta e_\phi &= c_G x + O(\|x\|^2), \qquad \Delta\bar{v} = c_v x + O(\|x\|^2), \end{aligned} \tag{F.17}$$

where $a = \Delta A_{\mathrm{exc}}$, $\Delta\Omega = \Omega - \Omega_*$ and $r_G = F_G(0, \Omega_*, A_*)$ is the nonzero base projection residual. Here Δe_ϕ is the change from the reference's fundamental phase error, not its absolute error relative to quadrature. Thus $A_G \in \mathbb{R}^{44\times 44}$, $p_G, q_G \in \mathbb{R}^{44}$, and $c_G, c_v \in \mathbb{R}^{1\times 44}$. The speed row c_v selects the constant v coefficient. If the reference's fundamental joint coefficients are $(a_{*,1}, b_{*,1})$, differentiating $\operatorname{atan2}(b_{*,1}, a_{*,1})$ gives the two nonzero phase-row entries $(-b_{*,1}, a_{*,1})/(a_{*,1}^2 + b_{*,1}^2)$.

The reported eigenvalues concern the homogeneous derivative model, with the constant residual and reference phase offset removed. Its phase channel approximates the rolling estimator by a first-order filter with bandwidth ω_c; h is the filtered phase perturbation and y the incremental integral contribution:

$$\dot{h} = \omega_c(c_G x - h), \qquad \dot{y} = K_I h, \qquad \Delta\Omega = s(K_P h + y). \tag{F.18}$$

Substituting $\Delta\Omega$ into the projected equation gives, for (x, h, y), the 46×46 phase-only matrix

$$\mathcal{A}_{\mathrm{PLL}} = \begin{bmatrix} A_G & sp_G K_P & sp_G \\ \omega_c c_G & -\omega_c & 0 \\ 0 & K_I & 0 \end{bmatrix}. \tag{F.19}$$

Equivalently, its phase-loop factor has the closed transfer form

$$\det(\lambda I - A_G)\left[\lambda(\lambda + \omega_c) - \omega_c s(K_P\lambda + K_I)c_G(\lambda I - A_G)^{-1}p_G\right] = 0. \tag{F.20}$$

The transfer expression is evaluated away from eigenvalues of A_G and extended as a characteristic polynomial after cancellation; its apparent poles are not additional roots.

At a fixed desired speed, the locally unsaturated perturbation of Equation (F.10) satisfies $\dot{a} = -K_v c_v x$ when $g_\phi = 1$ and no saturation is active. For the state (x, h, y, a), the 47×47 matrix used for the two-loop eigenvalues is

$$\mathcal{A}_{\mathrm{PLL+speed}} = \begin{bmatrix} A_G & sp_G K_P & sp_G & q_G \\ \omega_c c_G & -\omega_c & 0 & 0 \\ 0 & K_I & 0 & 0 \\ -K_v c_v & 0 & 0 & 0 \end{bmatrix}. \tag{F.21}$$

The numerical evaluation uses $s = +1$, $\omega_c = \Omega_{\mathrm{exc}}$, $K_P = 0.04$, $K_I = 0.08$, and $K_v = 2.5 \times 10^{-5}$. On the first fourteen ordered amplitude levels, $A_{\mathrm{exc}} = 0.0020496$ through 0.0314870, the spectral abscissa of $\mathcal{A}_{\mathrm{PLL}}$ lies in $[-1.81848, -1.00100] \times 10^{-3}$; for $\mathcal{A}_{\mathrm{PLL+speed}}$ it lies in $[-8.93731, -2.95865] \times 10^{-4}$. All 28 sampled spectral abscissae are negative. The next amplitude level, $A_{\mathrm{exc}} = 0.0355189$, is not included in this linearized set.

This result concerns the displayed finite-dimensional, unsaturated local models. The three-period speed-estimator window is replaced here by the instantaneous coefficient $c_v x$, so neither rolling estimator is reproduced exactly. The largest base projected-residual RMS is 2.44015×10^{-2}, and the smallest two-loop spectral margin is 2.95865×10^{-4}. These are different quantities, not directly comparable error bounds: no equilibrium-correction or perturbation bound transfers the matrix result to omitted harmonics, the sampled controller, or the exact nonlinear forced system.

F.7 What stronger convergence guarantees would require

The finite stage protocol can stop because it reaches its target, exhausts its proposal budget or encounters a rejected increment below the prescribed floor. Finite termination of that algorithm is therefore not a theorem of successful arrival at every requested speed.

Such a theorem would require an initial set inside a verified response neighborhood, uniform settling and acceptance for sufficiently small proposals, a positive minimum accepted progress step, and uniform stability and guard margins. Positive slopes at sampled responses do not establish a continuous lower slope bound, and negative eigenvalues of the truncated models do not supply these exact-plant assumptions. Even the sampled finite-harmonic evidence ends at amplitude 0.0314870, below the terminal controlled amplitudes near 0.03641. A stronger result must first justify its continuous operating interval and then cover the relevant endpoint. The present contribution remains the local phase–speed regulation observed in the specified simulations, supported by the stated finite-harmonic calculations.

Appendix G

Local geometry of continued locomotion families

The words *family*, *branch*, and *manifold* refer to related but different objects. This appendix makes their local mathematical relation explicit. It supports the terminology used in the abstract and conclusion; it also states why a finite numerical collection alone cannot establish one global smooth set.

G.1 Removing the arbitrary time shift

For an autonomous model, let z denote a pose-reduced mechanical state and let $\Phi_T(z_0; p)$ be the flow at time T for a fixed model parameter p. A periodic or reduced-periodic orbit has an arbitrary choice of time origin: if $z(t)$ is a solution, then $z(t+\tau)$ describes the same orbit. A section or phase condition removes this redundancy. In contrast, a periodically forced response sampled at a fixed forcing phase has already fixed its clock; adding an autonomous phase condition there would generally overconstrain the problem.

Collect the independent return equations and one phase condition in a residual

$$F(u) = 0, \qquad u = (z_0, T, \beta) \in \mathbb{R}^N, \qquad F : \mathbb{R}^N \to \mathbb{R}^M, \tag{G.1}$$

where β contains only genuine unknowns beyond the initial state and period: for example a model parameter, a forcing parameter, or an auxiliary unfolding variable. A derived quantity such as energy or mean speed can label the branch after solution. If it is promoted to an unknown, its defining equation must be included among the residual rows. The entries of F must be independent; in particular, a scalar exchange equation that replaces one return row must not also be counted as an additional row.

Proposition G.1 (Regular local branch)**.** *Suppose F is continuously differentiable near a solution $u_\star$ and $\operatorname{rank} DF(u_\star) = M$. Then the zero set of F is locally a C^1 manifold of dimension $N - M$. When $N - M = 1$, it is locally a continuation branch.*

Proof. This is the regular-level-set form of the implicit function theorem. A local choice of M dependent coordinates leaves $N - M$ free coordinates that parameterize the zero set. □

The proposition is conditional on the rank hypothesis. Solver convergence and a visually smooth plotted curve are not substitutes for that hypothesis. Pseudo-arclength continuation can traverse a simple fold of a chosen plotting parameter, but a branch point or other rank loss requires separate analysis (Kuznetsov, 2004).

G.2 From a phase-fixed branch to an invariant family

A phase-fixed branch at one fixed vector field selects one representative on every orbit. Restoring the phase sweeps each representative along its cycle:

$$\mathcal{M}_{\mathrm{loc}} = \{\Phi_\tau(z_0(\alpha); p) : \alpha \in I,\ \tau \in \mathbb{R}/T(\alpha)\mathbb{Z}\}\,. \tag{G.2}$$

For a local phase $\varphi = 2\pi\tau/T(\alpha)$, write the swept state as $Z(\alpha, \varphi)$. The two tangent vectors $\partial_\alpha Z$ and $\partial_\varphi Z$ must be independent; changing the cycle must not merely change its time origin. On a regular embedded

portion satisfying this rank-two condition, the sweep is locally two-dimensional: one coordinate selects the orbit and one selects phase on that orbit. It is invariant because advancing the fixed flow changes only the phase coordinate. If the branch coordinate changes the vector field itself, the branch instead lives in an extended orbit–parameter solution space; its union is not an invariant state-space manifold of any one member of that parameterized family.

This local statement does not imply that every computed branch, every modal sector, or different locomotor models belong to one global manifold. Branches may terminate, meet, bifurcate, or live in state spaces of different dimensions. The phrase *Natural Locomotion Manifold* names the organizing structure. The connected two-segment calculation supplies numerical transport and local derivative checks at 49 sampled nodes; the three-segment witnesses and atlas have their own distinct checks. Those finite computations support the stated local interpretations, not a global existence or regularity theorem.

G.3 When mean speed is a cycle coordinate

Let a be a regular coordinate along a fixed-mechanics branch and define $b(a) = T(a)^{-1} \int_0^{T(a)} v(t; a)\, \mathrm{d}t$. If $b'(a_\star) \neq 0$, the inverse function theorem gives a local inverse $a = a(b)$. Substitution in the phase sweep yields

$$z = Z\big(a(b), \varphi\big), \qquad b = \bar{v}. \tag{G.3}$$

It is a regular two-coordinate chart wherever the original sweep has rank two. A zero derivative of speed can invalidate this particular chart without destroying the underlying family. Numerical checks of a nonzero speed tangent at sampled nodes support its use there, not on every unseen branch. The passive dynamics changes phase, not b; moving between speed-labelled cycles requires an intervention.

G.4 A manifold and its observable image are not the same object

Let h collect reported quantities, for example mean speed, period, energy, a local phase coordinate, or pose increment:

$$h : \mathcal{M}_{\mathrm{loc}} \to \mathbb{R}^r. \tag{G.4}$$

Even if $\mathcal{M}_{\mathrm{loc}}$ is smooth, its image under h may fold or self-intersect. A fold in a speed–energy plot can therefore be a property of the chosen projection rather than a singularity of the underlying solution set. Conversely, a smooth-looking point cloud does not establish the rank, embeddedness, or global topology of its source.

Figure G.1 contrasts these two distinct phenomena: a regular source branch whose observable image folds, and a branch intersection where the regular-level-set argument fails. The upper example explains why a failed mean-speed coordinate may call for another branch parameter, not abandonment of the family. The lower example requires a separate analysis of the solution branches themselves.

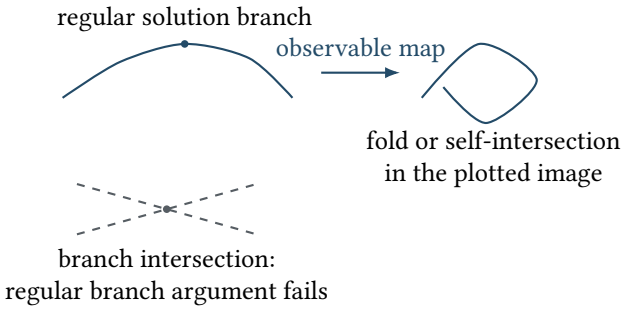


Figure G.1: Local solution geometry and an observable projection. A fold of a plotted coordinate need not destroy regularity of the source branch (top). The intersecting branches below cannot form a single regular one-dimensional chart at their intersection. This conceptual diagram contrasts a projection effect with a singular solution set; it is not a computed locomotion family.

G.5 Natural and resonant locomotion objects

For the conservative calculations, F contains exact return of the reduced state, a phase condition, and the fixed-model equations. Pose reconstruction then selects the transported subset $\Delta g \neq e_G$; the propulsive cycles reported here also have nonzero translation. Only regular embedded local phase sweeps of that subset are called NLM charts. In the two-segment scalar theorem, zero integrated exchange replaces the one remaining crossing speed equation after the other return coordinates have closed. In the three-segment locomotor, the independent vector-return conditions are solved and exchange is an endpoint identity or diagnostic; adding it as an independent selector would give a false codimension. Its two-dimensional equal-storage surface in proposition 5.1 is an endpoint level set, not the two-dimensional cycle–phase NLM.

For a periodically forced model, the state may be augmented by forcing phase, or equivalently sampled stroboscopically. For one-to-one synchronization, a residual of the form

$$F_{\mathrm{f}}(z_0, T_{\mathrm{exc}}, A_{\mathrm{exc}}, \ldots) = 0 \tag{G.5}$$

defines a synchronized-response solution set $\mathcal{S}_R$. If the independent residual Jacobian has the expected rank, regular portions form local curves or sheets according to the number of free forcing and design coordinates. Pose reconstruction then selects the transported subset

$$\mathcal{M}_R = \{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}. \tag{G.6}$$

Regular local components of this transported set, continued within the one-to-one resonance construction, are called the *Resonant Locomotion Manifold* (RLM). Here *resonant* identifies the response family organized around the one-to-one resonance; it does not require every member to lie at exact fundamental quadrature. Speed–phase–period diagrams are observable maps, not the definition of smoothness or transport.

In the forced two-segment construction used here, the stroboscopic unknowns are $(z_0, A, T) \in \mathbb{R}^6$ and there are four independent state periodicity equations. Full row rank therefore gives a local two-dimensional solution sheet at fixed forcing phase. Its coordinates select responses across forcing settings; they are not two independent mechanical phases. Restoring forcing phase yields a three-dimensional lifted object in state–phase–parameter space under the corresponding regularity assumptions. Parameters remain constant on each lifted trajectory. This differs from the two-dimensional cycle–phase NLM of one autonomous conservative field; the identical acronym suffix does not identify the spaces.

Pose reconstruction of all 13,782 recorded responses gives a nonidentity increment at the numerical resolution stated in Section E.5. Recorded closure, unit forcing-clock winding, and the population-wide exact-field local-flow, energy-work, and pose checks therefore support the numerical classification of the synchronized responses as transport-bearing, and hence as members of $\mathcal{M}_R$ at the stated resolution. The retained data do not include the complete collocation Jacobians needed for a population-wide regularity test. Calling a neighbourhood an RLM chart consequently remains conditional on the rank and embedding hypotheses. The mechanics-preserving local conservative-to-forced bridge is a separate calculation and does not prove persistence of the complete conservative family.

G.6 What a numerical regularity check establishes

Where the independent residual and its scaling are reproducible, a useful local check reports the singular values of the scaled Jacobian along a branch. The check must specify:

1. the unknowns and independent residual rows after phase fixing;
2. the scaling used for variables with different units or magnitudes;
3. the smallest singular value or numerical rank along regular segments;
4. the locations excluded as branch points, modal limits, or missing data.

The two-segment computation performs these local tests at its reported nodes; the three-segment and forced studies have the more limited coverage specified in their chapters. Such checks support local-chart language on the examined portions. They do not prove a single global NLM or RLM across different models, modal sectors, bifurcations, or parameterizations.

Bibliography

Gaëtan Abeloos, F. Müller, E. Ferhatoglu, Maren Scheel, Christophe Collette, Gaëtan Kerschen, Matthew R. W. Brake, Paolo Tiso, Ludovic Renson, and Malte Krack. A consistency analysis of phase-locked-loop testing and control-based continuation for a geometrically nonlinear frictional system. *Mechanical Systems and Signal Processing*, 170:108820, 2022. doi: 10.1016/j.ymssp.2022.108820.

Alin Albu-Schäffer and Cosimo Della Santina. A review on nonlinear modes in conservative mechanical systems. *Annual Reviews in Control*, 50:49–71, 2020. doi: 10.1016/j.arcontrol.2020.10.002.

Vladimir I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer, 2 edition, 1989. doi: 10.1007/978-1-4757-2063-1.

David A. W. Barton. Control-based continuation: Bifurcation and stability analysis for physical experiments. *Mechanical Systems and Signal Processing*, 84:54–64, 2017. doi: 10.1016/j.ymssp.2015.12.039.

Filip Bjelonic, Arne Sachtler, Cosimo Della Santina, and Alin Albu-Schäffer. Experimental closed-loop excitation of nonlinear normal modes on an elastic industrial robot. *IEEE Robotics and Automation Letters*, 7(2):1689–1696, 2022. doi: 10.1109/LRA.2022.3141156.

Anthony M. Bloch. *Nonholonomic Mechanics and Control*. Springer, 2 edition, 2015. doi: 10.1007/978-1-4939-3017-3.

Frédéric Boyer and Ayman Belkhiri. Reduced locomotion dynamics with passive internal DoFs: Application to nonholonomic and soft robotics. *IEEE Transactions on Robotics*, 30(3):578–592, 2014. doi: 10.1109/TRO.2013.2294733.

Bernard Brogliato. *Nonsmooth Mechanics: Models, Dynamics and Control*. Springer, 3 edition, 2016. doi: 10.1007/978-3-319-28664-8.

Davide Calzolari, Cosimo Della Santina, and Alin Albu-Schäffer. Exciting families of passive gaits in an elastic quadruped via natural motion manifold control. *The International Journal of Robotics Research*, page 02783649251347305, 2025. doi: 10.1177/02783649251347305.

Carlos Canudas-de Wit, Bernard Espiau, and Claudio Urrea. Orbital stabilization of underactuated mechanical systems. *IFAC Proceedings Volumes*, 35(1):527–532, 2002. doi: 10.3182/20020721-6-ES-1901.00900.

Mattia Cenedese and George Haller. How do conservative backbone curves perturb into forced responses? a Melnikov function analysis. *Proceedings of the Royal Society A*, 476(2234):20190494, 2020a. doi: 10.1098/rspa.2019.0494.

Mattia Cenedese and George Haller. Stability of forced–damped response in mechanical systems from a Melnikov analysis. *Chaos*, 30(8):083103, 2020b. doi: 10.1063/5.0012480.

Robert Clewley. Hybrid models and biological model reduction with PyDSTool. *PLOS Computational Biology*, 8(8):e1002628, 2012. doi: 10.1371/journal.pcbi.1002628.

Steven H. Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005. doi: 10.1126/science.1107799.

Alessandro De Luca, Giuseppe Oriolo, and Paolo Robuffo Giordano. Image-based visual servoing schemes for nonholonomic mobile manipulators. *Robotica*, 25(2):131–145, 2007. doi: 10.1017/S0263574706003262.

Cosimo Della Santina and Alin Albu-Schäffer. Exciting efficient oscillations in nonlinear mechanical systems through eigenmanifold stabilization. *IEEE Control Systems Letters*, 5(6):1916–1921, 2021. doi: 10.1109/LCSYS.2020.3048228.

Vivien Denis, Marc Jossic, Christophe Giraud-Audine, Benoît Chomette, Alexandre Renault, and Olivier Thomas. Identification of nonlinear modes using phase-locked-loop experimental continuation and normal form. *Mechanical Systems and Signal Processing*, 106:430–452, 2018. doi: 10.1016/j.ymssp.2018.01.014.

Eusebius J. Doedel, Alan R. Champneys, Fabio Dercole, Thomas F. Fairgrieve, Yuri A. Kuznetsov, Bart E. Oldeman, Randy C. Paffenroth, Björn Sandstede, Xianjun Wang, and Chenghai Zhang. *AUTO-07P: Continuation and Bifurcation Software for Ordinary Differential Equations*, 2007. URL https://auto-07p.github.io/.

Jaap Eldering and Henry O. Jacobs. The role of symmetry and dissipation in biolocomotion. *SIAM Journal on Applied Dynamical Systems*, 15(1):24–59, 2016. doi: 10.1137/140970914.

Adrien Escande, Nicolas Mansard, and Pierre-Brice Wieber. Hierarchical quadratic programming: Fast online humanoid-robot motion generation. *The International Journal of Robotics Research*, 33(7):1006–1028, 2014. doi: 10.1177/0278364914521306.

Francesco Fassò, Simone Passarella, and Marta Zoppello. Control of locomotion systems and dynamics in relative periodic orbits. *Journal of Geometric Mechanics*, 12(3):395–420, 2020. doi: 10.3934/jgm.2020022.

Vitaliy Fedonyuk and Phanindra Tallapragada. The dynamics of a chaplygin sleigh with an elastic internal rotor. *Regular and Chaotic Dynamics*, 24(1):114–126, 2019. doi: 10.1134/S1560354719010076.

Zhenyu Gan, Yevgeniy Yesilevskiy, Petr Zaytsev, and C. David Remy. All common bipedal gaits emerge from a single passive model. *Journal of the Royal Society Interface*, 15(146):20180455, 2018. doi: 10.1098/rsif.2018.0455.

Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Springer, 2 edition, 1993. doi: 10.1007/978-3-540-78862-1.

George Haller and Sten Ponsioen. Nonlinear normal modes and spectral submanifolds: Existence, uniqueness and use in model reduction. *Nonlinear Dynamics*, 86(3):1493–1534, 2016. doi: 10.1007/s11071-016-2974-z.

Ross L. Hatton and Howie Choset. Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice. *The International Journal of Robotics Research*, 30(8):988–1014, 2011. doi: 10.1177/0278364910394392.

Scott D. Kelly and Richard M. Murray. Geometric phases and robotic locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995. doi: 10.1002/rob.4620120607.

Gaëtan Kerschen, Maxime Peeters, Jean-Claude Golinval, and Alexander F. Vakakis. Nonlinear normal modes, part i: A useful framework for the structural dynamicist. *Mechanical Systems and Signal Processing*, 23(1):170–194, 2009. doi: 10.1016/j.ymssp.2008.04.002.

Robert J. Kuether, Ludovic Renson, Thibaut Detroux, Chiara Grappasonni, Gaëtan Kerschen, and Matthew S. Allen. Nonlinear normal modes, modal interactions and isolated resonance curves. *Journal of Sound and Vibration*, 351:299–310, 2015. doi: 10.1016/j.jsv.2015.04.035.

Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*. Springer, 3 edition, 2004. doi: 10.1007/978-1-4757-3978-7.

Nicolas Mansard, Oussama Khatib, and Abderrahmane Kheddar. A unified approach to integrate unilateral constraints in the stack of tasks. *IEEE Transactions on Robotics*, 25(3):670–685, 2009. doi: 10.1109/TRO.2009.2020345.

Tad McGeer. Passive dynamic walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990. doi: 10.1177/027836499000900206.

Mirado Mortel, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Natural locomotion: Principle and method, 2026. arXiv:2605.28254.

Jim P. Ostrowski and Joel W. Burdick. The geometric mechanics of undulatory robotic locomotion. *The International Journal of Robotics Research*, 17(7):683–701, 1998. doi: 10.1177/027836499801700701.

Maxime Peeters, Régis Viguié, Guillaume Sérandour, Gaëtan Kerschen, and Jean-Claude Golinval. Nonlinear normal modes, part ii: Toward a practical computation using numerical continuation techniques. *Mechanical Systems and Signal Processing*, 23(1):195–216, 2009. doi: 10.1016/j.ymssp.2008.04.003.

Maxime Peeters, Gaëtan Kerschen, and Jean-Claude Golinval. Dynamic testing of nonlinear vibrating structures using nonlinear normal modes. *Journal of Sound and Vibration*, 330(3):486–509, 2011. doi: 10.1016/j.jsv.2010.08.028.

Simon Peter and Remco I. Leine. Excitation power quantities in phase resonance testing of nonlinear systems with phase-locked-loop excitation. *Mechanical Systems and Signal Processing*, 96:139–158, 2017. doi: 10.1016/j.ymssp.2017.04.011.

Simon Peter, Robin Riethmüller, and Remco I. Leine. Tracking of backbone curves of nonlinear systems using phase-locked-loops. In *Nonlinear Dynamics, Volume 1*, pages 107–120. Springer, 2016. doi: 10.1007/978-3-319-29739-2_11.

Hari Krishna Hari Prasad, Ross L. Hatton, and Kaushik Jayaram. Geometric mechanics of contact-switching systems. *IEEE Robotics and Automation Letters*, 8(12):8343–8349, 2023. doi: 10.1109/LRA.2023.3327930.

Maximilian Raff, Nelson Rosa, Jr., and C. David Remy. Connecting gaits in energetically conservative legged systems. *IEEE Robotics and Automation Letters*, 7(3):8407–8414, 2022. doi: 10.1109/LRA.2022.3186500.

Mirado Rajaomarosata, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Natural efficient gaits from nonholonomic locomotion nonlinear normal mode (NL-NNM): The pendrivencar case. *Mechatronics*, 110:103366, 2025a. doi: 10.1016/j.mechatronics.2025.103366.

Mirado Rajaomarosata, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Energy-efficient nonholonomic fish robot: Nonlinear forced oscillations. In *OCEANS 2025 Brest*, pages 1–7. IEEE, 2025b. doi: 10.1109/OCEANS58557.2025.11104676.

Sophie Ramananarivo, Ramiro Godoy-Diana, and Benjamin Thiria. Rather than resonance, flapping wing flyers may play on aerodynamics to improve performance. *Proceedings of the National Academy of Sciences*, 108(15):5964–5969, 2011. doi: 10.1073/pnas.1017910108.

C. David Remy. *Optimal Exploitation of Natural Dynamics in Legged Locomotion*. PhD thesis, ETH Zurich, 2011.

Ludovic Renson, Alicia Gonzalez-Buelga, David A. W. Barton, and Simon A. Neild. Robust identification of backbone curves using control-based continuation. *Journal of Sound and Vibration*, 367:145–158, 2016. doi: 10.1016/j.jsv.2015.12.035.

R. M. Rosenberg. On nonlinear vibrations of systems with many degrees of freedom. In *Advances in Applied Mechanics*, volume 9, pages 155–242. Elsevier, 1966. doi: 10.1016/S0065-2156(08)70008-5.

Layale Saab, Oscar E. Ramos, François Keith, Nicolas Mansard, Philippe Souères, and Jean-Yves Fourquet. Dynamic whole-body motion generation under rigid contacts and other unilateral constraints. *IEEE Transactions on Robotics*, 29 (2):346–362, 2013. doi: 10.1109/TRO.2012.2234351.

Maren Scheel, Simon Peter, Remco I. Leine, and Malte Krack. A phase resonance approach for modal testing of structures with nonlinear dissipation. *Journal of Sound and Vibration*, 435:56–73, 2018. doi: 10.1016/j.jsv.2018.07.010.

Johan K. Schönebaum, Farbod Alijani, and Giuseppe Radaelli. Review on mobile robots that exploit resonance. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 235(24):7907–7924, 2021. doi: 10.1177/09544062211036923.

Steven W. Shaw and Christophe Pierre. Normal modes for non-linear vibratory systems. *Journal of Sound and Vibration*, 164(1):85–124, 1993. doi: 10.1006/jsvi.1993.1198.

Martin Tang, Cyrille Stéphan, and Marc Böswald. Phase resonance method for nonlinear mechanical structures with phase locked loop control. In *Proceedings of ISMA 2020*, pages 1805–1818, 2020.

Eric D. Tytell, Chia-Yu Hsu, and Lisa J. Fauci. The role of mechanical resonance in the neural control of swimming in fishes. *Zoology*, 117(1):48–56, 2014. doi: 10.1016/j.zool.2013.10.011.

Eric R. Westervelt, Jessy W. Grizzle, and Daniel E. Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1):42–56, 2003. doi: 10.1109/TAC.2002.806653.

COLLEGES SCIENCES
BRETAGNE POUR L'INGENIEUR
LOIRE ET LE NUMERIQUE

ENSTA INSTITUT POLYTECHNIQUE DE PARIS

Titre : Familles de locomotion naturelle : principe, méthode et commande

Mots clés : locomotion naturelle ; modes non linéaires ; mécanique géométrique ; échange d'énergie ; résonance ; commande

Résumé : Quels mouvements la mécanique d'un robot favorise-t-elle, et comment les entretenir sans imposer une trajectoire articulaire ? Cette thèse aborde la question dans un modèle idéal sans dissipation ni actionnement, puis réintroduit les pertes et l'apport d'énergie.

La locomotion naturelle répète l'état interne du robot tandis que son corps se déplace. Un principe relie cette répétition au bilan d'énergie entre un canal oscillatoire effectif et un canal propulsif, couplés par les contraintes avec l'environnement. Pour un locomoteur à deux segments, un théorème établit une équivalence entre ce bilan et la répétition de la vitesse articulaire sur une section orientée, sous les autres conditions de périodicité et de déplacement. La méthode calcule des familles de cycles, décrites localement par une variété de locomotion naturelle. Avec trois segments, elle recherche la périodicité de tout l'état interne et distingue des coordinations en phase et en opposition de phase. Les calculs éclairent aussi le choix des paramètres mécaniques.

Avec dissipation, une continuation relie localement un cycle naturel à une réponse entretenue. Une étude forcée distincte organise des familles de locomotion résonnante. Une boucle à verrouillage de phase ajuste la fréquence du couple, tandis qu'une boucle plus lente règle son amplitude pour maintenir la vitesse moyenne. Les résultats sont analytiques et numériques : l'actionnement accompagne un mouvement dont la mécanique détermine la forme.

Title: Natural Locomotion Families: Principle, Method, and Control

Keywords: natural locomotion; nonlinear modes; geometric mechanics; energy exchange; resonance; control

Abstract: Which motions does a robot's mechanics favour, and how can they be maintained without prescribing a joint trajectory? This thesis first studies an ideal model without dissipation or actuation, then reintroduces losses and energy input.

Natural locomotion repeats the robot's internal state while its body moves. A principle connects this repetition to the energy balance between an effective oscillatory channel and a propulsive channel, coupled through environmental constraints. For a two-segment locomotor, a theorem establishes an equivalence between this balance and repetition of the joint rate on an oriented section, subject to the remaining periodicity and displacement conditions. The method computes cycle families locally described by a Natural Locomotion Manifold. With three segments, it seeks periodicity of the entire internal state and distinguishes in-phase and anti-phase organizations. The computations also inform mechanical-parameter choices.

With dissipation, continuation locally connects a natural cycle to a maintained response. A separate forced study organizes resonant locomotion families. A phase-locked loop adjusts torque frequency, while a slower loop adjusts its amplitude to regulate mean speed. The results are analytical and numerical: actuation accompanies a motion whose shape is determined by the mechanics.

Lab-STICC – UMR CNRS 6285 – équipe ROBEX
ENSTA, 2 rue François Verny, 29806 Brest Cedex 09, France.